arXiv:1603.06916v2 [math.OC] 19 May 2016

SOLVING GENERIC NONARCHIMEDEAN SEMIDEFINITE PROGRAMS USING STOCHASTIC GAME ALGORITHMS

XAVIER ALLAMIGEON, STÉPHANE GAUBERT, AND MATEUSZ SKOMRA

Abstract. A general issue in computational optimization is to develop combinatorial algorithms for semidefinite programming. We address this issue when the base field is *nonarchimedean*. We provide a solution for a class of semidefinite feasibility problems given by generic matrices with a Metzler-type sign pattern. Our approach is based on tropical geometry. We define tropical spectrahedra as the images by the valuation of nonarchimedean spectrahedra, and provide an explicit description of the tropical spectrahedra arising from the aforementioned class of problems. We deduce that the tropical semidefinite feasibility problems obtained in this way are equivalent to stochastic mean payoff games, which have been well studied in algorithmic game theory. This allows us to solve nonarchimedean semidefinite feasibility problems using algorithms for stochastic games. These algorithms are of a combinatorial nature and work for large instances.

1. Introduction

Semidefinite programming consists in optimizing a linear function over a spectrahedron. The latter is a subset of $\mathbb{R}^n$ defined by linear matrix inequalities, i.e., a set of the form

$$\mathcal{S} = \{x \in \mathbb{R}^n \colon Q^{(0)} + x_1 Q^{(1)} + \cdots + x_n Q^{(n)} \succcurlyeq 0\}$$

where the $Q^{(k)}$ are symmetric matrices of order m, and $\succcurlyeq$ denotes the Loewner order on the space of symmetric matrices. By definition, $X \succcurlyeq Y$ if and only if $X - Y$ is positive semidefinite.

Semidefinite programming is a fundamental tool in convex optimization. It is used to solve various applications from engineering sciences, and also to obtain approximate solutions or bounds for hard problems arising in combinatorial optimization and semialgebraic optimization. We refer the reader to [BPR13, GM12] for information.

Semidefinite programs are usually solved via interior point methods. The latter provide an approximate solution in a polynomial number of iterations, provided that a strictly feasible initial solution, i.e., a point belonging to the interior of the set $\mathcal{S}$, is known. We refer the reader to [dKV15] for a detailed analysis of the complexity in the bit model of interior point methods for semidefinite programming, and for a discussion of earlier complexity results based on the ellipsoid method.

Semidefinite programming becomes a much harder matter if one requires an exact solution. The feasibility problem (deciding the emptiness of the set $\mathcal{S}$) belongs to $\mathrm{NP}_{\mathbb{R}} \cap \mathrm{coNP}_{\mathbb{R}}$, where the subscript $\mathbb{R}$ refers to the BSS model of computation [Ram97]. It is not known to be in NP in the bit model. A difficulty here is that all feasible points may have entries of absolute value doubly exponential in the size of the input. Also, there may be no rational solution [Sch13]. Beyond their theoretical interest, exact algorithms for semidefinite programming may be useful to address problems of formal proofs, which sometimes lead to challenging (degenerate) instances [MC11]. Known exact methods rely either on general purpose semialgebraic techniques (like cylindrical decomposition) or on dedicated methods based on the computation of critical points, see the recent work [HNSED15] and the references therein.

Date: May 20, 2016.

The three authors were partially supported by the ANR projects CAFEIN (ANR-12-INSE-0007) and MALTHY (ANR-13-INSE-0003), and by the PGMO program of EDF and Fondation Mathématique Jacques Hadamard. M. Skomra is supported by a grant from Région Ile-de-France.

Semidefinite programing is meaningful in any real closed field, and in particular in *non-archimedean* (real closed) fields like the field of Puiseux series with real coefficients. Such nonarchimedean semidefinite programming problems arise when considering *parametric* semidefinite programming problems over the reals, or structured problems in which the entries of the matrices have different orders of magnitudes. They are also of an intrinsic interest, since, by analogy with the situation in linear programming [Meg89], shifting to the nonarchimedean case is expected to shed light on the complexity of the classical problem over the reals. We note that the fields of formal Puiseux series are representative of the general nonarchimedean situation, since any nonarchimedean ordered field can be embedded in a field of generalized (Hahn) formal series [Ste10, Th. 5.2.20].

Description of main results. In this paper, we address semidefinite programming in the *non-archimedean* case, to which methods from tropical geometry can be applied. These methods are expected to allow one, in generic situations, to reduce semialgebraic problems to combinatorial problems, involving only the nonarchimedean valuations (leading exponents) of the coefficients of the input. We exploit the tropical approach by considering tropical spectrahedra, defined as the images by the valuation of nonarchimedean spectrahedra.

We first give a combinatorial description (Theorem 3.4 and Theorem 3.7) of tropical spectrahedra arising as the images of generic nonarchimedean spectrahedra determined by symmetric matrices that satisfy a Metzler-type sign condition (nonpositivity of off-diagonal entries). The subclass of inputs under consideration already covers known hard examples of semidefinite feasibility problems, involving for instance cascades of quadratic inequalities, as in [Ram97]. The genericity assumption involves valuations. Assumptions of this nature are familiar in tropical geometry.

Then, we show (Corollaries 4.5 and 4.6) that the tropical semidefinite feasibility problem for the aforementioned class is equivalent to solving stochastic mean payoff games. The latter are a well-studied class of combinatorial zero-sum games with perfect information. The complexity of these games is a long-standing open problem. They belong to the class $\mathrm{NP} \cap \mathrm{coNP}$, but they are not known to be polynomial.

This allows us to apply stochastic game algorithms to solve nonarchimedean semidefinite feasibility problems (Section 4.3). We obtain in this way both theoretical bounds and a practicable method which solves some large scale instances. We point out that this paper is an extended version of the proceedings article [AGS16].

Related work. It is a general principle in tropical geometry that some information on the solutions of polynomial systems can be obtained by combinatorial reasoning about valuations. Early incarnations of this idea are the patchworking method of Viro to construct real curves with a prescribed topology [Vir89] and its extension by Sturmfels [Stu94] to the case of complete intersections of hypersurfaces. See also the more recent work by Bihan [Bih02]. Our characterization of tropical spectrahedra is obtained by applying the same kind of ideas to a class of *semialgebraic* systems which arise from semidefinite feasibility problems (we deal with inequalities, instead of equalities).

Yu [Yu15] showed that the image by the valuation of the cone of nonarchimedean positive semidefinite matrices is described only by tropical minors of order 2. Our combinatorial characterization of tropical Metzler spectrahedra can be thought of as an extension of her theorem to the case of nonarchimedean spectrahedra.

This work also originates from the equivalence between tropical polyhedral feasibility problems and deterministic zero-sum games with mean payoff, established by Akian, Gaubert, and Guterman [AGG12]. The novelty here is the handling of nonlinear semialgebraic convex problems, and the proof that the well-known class of *stochastic* mean payoff games correspond precisely to semidefinite feasibility problems with a Metzler structure, therefore relating two classes of problems which both have an unsettled complexity.

Finally, as mentioned above, the computation of exact solutions of semidefinite programming problems is of current interest. In particular, Nie, Ranestad, and Sturmfels [NRS10] provided

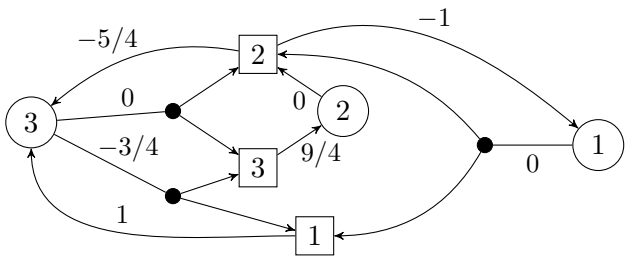

FIGURE 1. Stochastic game associated with the spectrahedron considered in Section 2.

complexity measures based on the notion of algebraic degree, and dedicated algorithms have been developed by Henrion, Naldi, and Safey El Din [HNSED15].

2. Statement of the problem and illustration of our approach

The nonarchimedean structure which we use in this paper is the *field* $\mathbb{K}$ *of (formal generalized real) Puiseux series*, which is composed of formal series of the form

$$\boldsymbol{x} = \sum_{i=1}^{\infty} c_{\lambda_i} t^{\lambda_i}\,, \tag{1}$$

where t is a formal parameter, $(\lambda_i)_{i\geqslant 1}$ is a strictly decreasing sequence of real numbers that is either finite or unbounded, and $c_{\lambda_i} \in \mathbb{R} \setminus \{0\}$ for all λ_i. There is also a special, empty series, which is denoted by 0. We denote by $\mathsf{lc}(\boldsymbol{x})$ the coefficient c_{λ_1} of the leading term in the series $\boldsymbol{x}$, with the convention that $\mathsf{lc}(0) = 0$. The addition and multiplication in $\mathbb{K}$ are defined in a natural way. Moreover, $\mathbb{K}$ can be endowed with a linear order $\geqslant$, which is defined as $\boldsymbol{x} \geqslant \boldsymbol{y}$ if $\mathsf{lc}(\boldsymbol{x} - \boldsymbol{y}) \geqslant 0$. We denote $\mathbb{K}_{\geqslant 0}$ the set of nonnegative series $\boldsymbol{x}$, i.e., satisfying $\boldsymbol{x} \geqslant 0$. The *valuation* of an element $\boldsymbol{x} \in \mathbb{K}$ as in (1) is defined as the greatest exponent λ_1 occurring in the series. It is known that $\mathbb{K}$ is a real closed field (see [Mar10] for instance).

In this paper, we consider a homogeneous variant of the semidefinite feasibility problem over the field $\mathbb{K}$. More precisely, given symmetric matrices $\boldsymbol{Q}^{(1)}, \dots, \boldsymbol{Q}^{(n)} \in \mathbb{K}^{m\times m}$, the problem which we study consists in determining whether or not the following spectrahedral cone

$$\{\boldsymbol{x} \in \mathbb{K}^n_{\geqslant 0} \colon \boldsymbol{x}_1\boldsymbol{Q}_1 + \cdots + \boldsymbol{x}_n\boldsymbol{Q}_n \succcurlyeq 0\}$$

is *trivial*, meaning that it is reduced to the identically null point of $\mathbb{K}^n_{\geqslant 0}$. The linear map $\boldsymbol{x} \mapsto \boldsymbol{x}_1\boldsymbol{Q}_1 + \cdots + \boldsymbol{x}_n\boldsymbol{Q}_n$ is a *matrix pencil*, which we denote by $\boldsymbol{Q}(\boldsymbol{x})$ for more brevity. The latter decision problem is in fact equivalent to the nontriviality problem of general spectrahedral cones in which the nonnegativity condition $\boldsymbol{x} \in \mathbb{K}^n_{\geqslant 0}$ is relaxed. Indeed, given a pencil $\boldsymbol{Q}(\boldsymbol{x})$ of symmetric matrices, the spectrahedron $\{\boldsymbol{x} \in \mathbb{K}^n \colon \boldsymbol{Q}(\boldsymbol{x}) \succcurlyeq 0\}$ is trivial if and only if the spectrahedron $\{(\boldsymbol{y}, \boldsymbol{z}) \in \mathbb{K}^{2n}_{\geqslant 0} \colon (\boldsymbol{y}_1 - \boldsymbol{z}_1)\boldsymbol{Q}_1 + \cdots + (\boldsymbol{y}_n - \boldsymbol{z}_n)\boldsymbol{Q}_n \succcurlyeq 0\}$ is trivial. Even if the instances arising in this way are unlikely to be generic in the sense we define later in Section 3, we consider that the decision problem which we focus on already retains much of the complexity of the semidefinite feasibility problem over the field of Puiseux series.

Our approach applies to the case where the off-diagonal entries of the matrices $\boldsymbol{Q}^{(1)}$, …, $\boldsymbol{Q}^{(n)}$ are nonpositive. Then, we associate to these matrices the following zero-sum game. There are two players, Player Min and Player Max, who control disjoint sets of states. The states of Player Min can be identified to the variables $\boldsymbol{x}_1, \dots, \boldsymbol{x}_n$. The states of Player Max can be identified to the row (or column) indices of the matrices, i.e., to the elements of $\{1, \dots, m\}$. In state $\boldsymbol{x}_k$, Player Min choses a subset $\{i, j\} \subset \{1, \dots, m\}$ such that the entry $\boldsymbol{Q}^{(k)}_{ij}$ is negative. Next, Player Min pays to Player Max the opposite of the valuation of $\boldsymbol{Q}^{(k)}_{ij}$. Then, Nature selects one element among i or j at random, with uniform probabilities (i.e., i or j is drawn

with probability 1/2).[1] If i is drawn, meaning that the current state is now i, Player Max chooses a variable $\boldsymbol{x}_l$ such that $\boldsymbol{Q}_{ii}^{(l)}$ has a positive sign. He receives the valuation of $\boldsymbol{Q}_{ii}^{(l)}$ from Player Min, and the next state becomes $\boldsymbol{x}_l$. If j is drawn, the rule of move and the payment are identical, up to the replacement of i by j. A (stationary) policy of one player is a map which associates to a state of the player an admissible move. We are interested in infinite plays, with an infinite number of turns. Player Max looks for a policy which maximizes the mean payment received from Player Min per time unit, while Player Min looks for a policy which minimizes it.

We can think informally of this construction as follows: Player Min wishes to show that the semidefinite programming problem is infeasible, whereas Player Max wishes to show that it is feasible.

Let us now illustrate our approach on an example in dimension 3. We consider the following pencil of symmetric matrices

$$\boldsymbol{Q}(\boldsymbol{x}) := \begin{bmatrix} t\boldsymbol{x}_3 & -\boldsymbol{x}_1 & -t^{3/4}\boldsymbol{x}_3 \\ -\boldsymbol{x}_1 & t^{-1}\boldsymbol{x}_1 + t^{-5/4}\boldsymbol{x}_3 - \boldsymbol{x}_2 & -\boldsymbol{x}_3 \\ -t^{3/4}\boldsymbol{x}_3 & -\boldsymbol{x}_3 & t^{9/4}\boldsymbol{x}_2 \end{bmatrix},$$

and we aim at checking whether or not the spectrahedron $\{\boldsymbol{x} \in \mathbb{K}^3_{\geqslant 0} \colon \boldsymbol{Q}(\boldsymbol{x}) \succcurlyeq 0\}$ is trivial. The associated game is depicted on Figure 1. The states of Player Min are depicted by circles. The states of Player Max are depicted by squares. The states in which Nature plays are depicted by full dots. The admissible moves of the game are represented by the edges between the states. The corresponding payments received by Player Max are indicated on these edges.

Observe that both players in this game have only two policies: at state $\boxed{2}$ Player Max can choose the move that goes to ① or the one that goes to ③, whereas at state ③ Player Min can choose the move $\{\boxed{1}, \boxed{3}\}$ or the move $\{\boxed{2}, \boxed{3}\}$.

Suppose that Player Max chooses the policy which goes to state ① from state $\boxed{2}$. If Player Min plays using the move $\{\boxed{1}, \boxed{3}\}$ at state ③, then a standard computation on Markov chain (which we present in Appendix A.2) shows that the long-term average payoff of Player Max is equal to 3/40. Similarly, if Player Min chooses the move $\{\boxed{2}, \boxed{3}\}$ instead, then the payoff of Player Max is equal to 1/56. Therefore, Player Max has a policy which guarantees him to win the game, i.e., to obtain a nonnegative payoff. The main theorem of this paper states that this fact is equivalent to the nontriviality of the spectrahedron $\{\boldsymbol{x} \in \mathbb{K}^3_{\geqslant 0} \colon \boldsymbol{Q}(\boldsymbol{x}) \succcurlyeq 0\}$. As we shall see, a winning policy of Player Max (resp. Min) provides a feasibility (resp. infeasibility) certificate. The mean payoff represents a feasibility/infeasibility margin.

3. Tropical spectrahedra

3.1. **Tropical algebra.** In this section, we briefly introduce the basic concepts of tropical algebra and its connection with the nonarchimedean field of Puiseux series.

We denote by $\mathsf{val} \colon \mathbb{K} \to \mathbb{R} \cup \{-\infty\}$ the map which associates a Puiseux series $\boldsymbol{x} \in \mathbb{K}$ to its valuation. We use the convention $\mathsf{val}(0) = -\infty$. It is immediate to see that the map val satisfies the following properties

$$\mathsf{val}(\boldsymbol{x} + \boldsymbol{y}) \leqslant \max(\mathsf{val}(\boldsymbol{x}), \mathsf{val}(\boldsymbol{y})) \tag{2}$$

$$\mathsf{val}(\boldsymbol{x}\boldsymbol{y}) = \mathsf{val}(\boldsymbol{x}) + \mathsf{val}(\boldsymbol{y}) \tag{3}$$

meaning that val is a *nonarchimedean valuation*. Moreover, the equality holds in (2) if the leading terms of $\boldsymbol{x}$ and $\boldsymbol{y}$ do not cancel, which is the case if $\mathsf{val}(\boldsymbol{x}) \neq \mathsf{val}(\boldsymbol{y})$ or if $\boldsymbol{x}, \boldsymbol{y} \geqslant 0$.

Loosely speaking, the *tropical semifield* $\mathbb{T}$ can be thought of as the image of $\mathbb{K}$ by the nonarchimedean valuation. The base set of $\mathbb{T}$ is defined to be $\mathbb{R} \cup \{-\infty\}$. It is endowed with the addition $x \oplus y := \max(x, y)$ and the multiplication $x \odot y := x + y$. The term semifield refers to the fact that the addition does not have an opposite law. We use the notation $\bigoplus_{i=1}^n a_i = a_1 \oplus \ldots \oplus a_n$ and $a^{\odot n} = a \odot \ldots \odot a$ (n times). We also endow $\mathbb{T}$ with the standard order $\leqslant$. The map val

[1] We allow the case $i = j$. In other words, Player Min can choose a subset $\{i\} \subset \{1, \ldots, m\}$ such that $\boldsymbol{Q}_{ii}^{(k)}$ is negative. In this case, Nature selects i.

yields an order-preserving morphism of semifields from $\mathbb{K}_{\geqslant 0}$ to $\mathbb{T}$. This follows from (3) and from the equality case in (2). We refer the reader to [But10, MS15] for more information on the tropical semifield.

It is convenient to keep track not only of the valuation of a series but also of its sign. To this end, we introduce the *signed valuation*, denoted by sval, which associates to a series $\boldsymbol{x} \in \mathbb{K}$ the couple $(\mathsf{sign}(\boldsymbol{x}), \mathsf{val}(\boldsymbol{x}))$, where $\mathsf{sign}(\boldsymbol{x}) \in \{0, \pm 1\}$ denotes the *sign* of $\boldsymbol{x}$, which is equal to $+1$ if $\boldsymbol{x} > 0$, -1 if $\boldsymbol{x} < 0$, and 0 otherwise. We denote by $\mathbb{T}_\pm$ the image of $\mathbb{K}$ by sval. We call it the set of *signed tropical numbers*. For brevity, we denote an element of the form (ϵ, a) by a if $\epsilon = 1$, $\ominus a$ if $\epsilon = -1$, and $-\infty$ if $\epsilon = 0$. Here, $\ominus$ is a formal symbol. We call the elements of the first and second kind the *positive* and *negative* tropical numbers, respectively. We denote by $\mathbb{T}_+$ and $\mathbb{T}_-$ the corresponding sets. In this way, (-2) is tropically positive, but $\ominus(-2)$ is tropically negative. Also, $\mathbb{T}$ is embedded in $\mathbb{T}_\pm$, i.e., $\mathbb{T} = \mathbb{T}_+ \cup \{-\infty\}$. We shall extend the valuation maps val and sval to vectors and matrices in a coordinate-wise manner.

In $\mathbb{T}_\pm$, we define a *modulus* function, $|\cdot| \colon \mathbb{T}_\pm \to \mathbb{T}$, as $|-\infty| = -\infty$ and $|a| = |\ominus a| = a$ for all $a \in \mathbb{T}_+$. We point out that $\odot$ straightforwardly extends to $\mathbb{T}_\pm$ using the standard rules for the sign, for instance $2 \odot (\ominus 3) = \ominus 5$. In contrast, we only partially extend the tropical addition $\oplus$ to elements of $\mathbb{T}_\pm$ of identical sign, e.g., $2 \oplus 3 = 3$ and $(\ominus 2) \oplus (\ominus 3) = \ominus 3$.

Finally, we use the notion of tropical polynomials. A *tropical (signed) polynomial* over the variables $X_1, \dots, X_n$ is a formal expression of the form

$$P(X) = \bigoplus_{\alpha \in \Lambda} a_\alpha \odot X_1^{\odot \alpha_1} \odot \dots \odot X_n^{\odot \alpha_n}\,, \tag{4}$$

where $\Lambda \subset \{0, 1, 2, \dots\}^n$, and $a_\alpha \in \mathbb{T}_\pm \setminus \{-\infty\}$ for all $\alpha \in \Lambda$. We say that the tropical polynomial P *vanishes* on the point $x \in \mathbb{T}_\pm^n$ if the terms $a_\alpha \odot x_1^{\odot \alpha_1} \odot \dots \odot x_n^{\odot \alpha_n}$ which have the greatest modulus do not have the same sign. If P does not vanish on x, we define $P(x)$ as the tropical sum of the terms which have the greatest modulus. As an example, if $P(X) = 2 \odot X_1^{\odot 3} \odot X_2^{\odot 4} \oplus (\ominus 0) X_2$, then $P(1, \ominus 5) = 25$, $P(1, -5) = \ominus(-5)$, whereas P vanishes on $(1, -5/3)$. These definitions are motivated by the following immediate lemma, which shows that the structure laws of $\mathbb{T}_\pm$ are essentially the images of the ones of $\mathbb{K}$.

Lemma 3.1. *Let*

$$\boldsymbol{P}(X) = \sum_{\alpha \in \Lambda} \boldsymbol{a}_\alpha X_1^{\alpha_1} \dots X_n^{\alpha_n} \in \mathbb{K}[X_1, \dots, X_n]$$

and let P be defined as in (4) with $a_\alpha := \mathsf{sval}(\boldsymbol{a}_\alpha)$. Then, for all $\boldsymbol{x} \in \mathbb{K}^n$,

$$\mathsf{sval}(\boldsymbol{P}(\boldsymbol{x})) = P(\mathsf{sval}(\boldsymbol{x}))\,,$$

provided that P does not vanish on $\mathsf{sval}(\boldsymbol{x})$. □

Given a polynomial $\boldsymbol{P}$ as in Lemma 3.1, we denote by $\boldsymbol{P}^+$ the polynomial formed by the terms $\boldsymbol{a}_\alpha X_1^{\alpha_1} \dots X_n^{\alpha_n}$ such that $\boldsymbol{a}_\alpha > 0$. Similarly, $\boldsymbol{P}^-$ refers to the polynomial consisting of the terms $-\boldsymbol{a}_\alpha X_1^{\alpha_1} \dots X_n^{\alpha_n}$ verifying $\boldsymbol{a}_\alpha < 0$. In this way, $\boldsymbol{P} = \boldsymbol{P}^+ - \boldsymbol{P}^-$. We also use the analogues of these polynomials in the tropical setting. If P is the tropical polynomial given in (4), we define P^+ (resp. P^-) as the tropical polynomial generated by the terms $|a_\alpha| \odot X_1^{\odot \alpha_1} \odot \dots \odot X_n^{\odot \alpha_n}$ where $a_\alpha \in \mathbb{T}_+$ (resp. $\mathbb{T}_-$). Observe that the quantities $P^+(x)$ and $P^-(x)$ are well defined for all $x \in \mathbb{T}^n$, since the tropical polynomials P^+ and P^- only involve tropically positive coefficients.

Throughout the paper, we denote the set $\{1, \dots, k\}$ by $[k]$.

3.2. **Tropicalization of nonarchimedean spectrahedra.** We now introduce the notion of tropical spectrahedra.

Definition 3.2. A set $\mathcal{S} \subset \mathbb{T}^n$ is said to be a *tropical spectrahedron* if there exists a spectrahedron $\boldsymbol{\mathcal{S}} \subset \mathbb{K}^n_{\geqslant 0}$ such that $\mathcal{S} = \mathsf{val}(\boldsymbol{\mathcal{S}})$.

If $\mathcal{S} = \mathsf{val}(\boldsymbol{\mathcal{S}})$, then we refer to $\mathcal{S}$ as the *tropicalization* of the spectrahedron $\boldsymbol{\mathcal{S}}$, and $\boldsymbol{\mathcal{S}}$ is said to be a *lift (over the field $\mathbb{K}$)* of $\mathcal{S}$. Our approach to the problem of checking the triviality of a spectrahedron $\boldsymbol{\mathcal{S}} \subset \mathbb{K}^n_{\geqslant 0}$ consists in determining whether or not the corresponding tropical

spectrahedron $\mathcal{S} = \mathsf{val}(\boldsymbol{\mathcal{S}})$ is *trivial*, i.e., is reduced to the singleton $\{(-\infty, \dots, -\infty)\}$. To this purpose, it is convenient to exploit some explicit description of the set $\mathcal{S}$. Recall that a symmetric matrix with entries in a real closed field is positive semidefinite if and only if every principal minor of the matrix is nonnegative. This property provides a description of $\boldsymbol{\mathcal{S}}$ by a system of polynomial inequalities of the form $\det \boldsymbol{Q}_{I\times I}(\boldsymbol{x}) \geqslant 0$, where I is a nonempty subset of $[m]$, and $\det \boldsymbol{Q}_{I\times I}(\boldsymbol{x})$ corresponds to the $(I \times I)$-minor of the matrix $\boldsymbol{Q}(\boldsymbol{x}) = \boldsymbol{x}_1\boldsymbol{Q}^{(1)} + \cdots + \boldsymbol{x}_n\boldsymbol{Q}^{(n)}$. Following this, we obtain that the tropical spectrahedron $\mathcal{S}$ is included in the intersection of the images of the sets $\{\boldsymbol{x} \in \mathbb{K}^n_{\geqslant 0} \colon \det \boldsymbol{Q}_{I\times I}(\boldsymbol{x}) \geqslant 0\}$ under the valuation map. Unfortunately, this inclusion may be strict, see [ABGJ15, Example 15] for an example in which $\boldsymbol{\mathcal{S}}$ is a polyhedron.

Section 3.2.1 deals with a class of tropical spectrahedra for which we overcome this difficulty, and obtain an explicit description. It turns out that this description only involves the tropical counterpart of the nonnegativity constraints on minors of order 1 and 2. In Section 3.2.2, we show that this kind of description is valid for the tropicalization of a large class of spectrahedra, for which the valuation of the coefficients of the matrices $\boldsymbol{Q}^{(k)}$ satisfy a genericity condition. Section 3.2.3 provides an illustration of these results on the spectrahedron introduced in Section 2.

3.2.1. *Tropical Metzler spectrahedra.* Recall that a matrix A is called *(negated) Metzler matrix* if its off-diagonal coefficients are nonpositive. Similarly, we say that a matrix $M \in \mathbb{T}^{m\times m}_{\pm}$ is a *tropical Metzler matrix* if $M_{ij} \in \mathbb{T}_- \cup \{-\infty\}$ for all $i \neq j$. Let $Q^{(1)}, \dots, Q^{(n)} \in \mathbb{T}^{m\times m}_{\pm}$ be symmetric tropical Metzler matrices. Given $i, j \in [m]$, we refer to $Q_{ij}(X)$ as the tropical polynomial:

$$Q_{ij}(X) := Q^{(1)}_{ij} \odot X_1 \oplus \ldots \oplus Q^{(n)}_{ij} \odot X_n \,.$$

Definition 3.3. If $Q^{(1)}, \dots, Q^{(n)} \in \mathbb{T}^{m\times m}_{\pm}$ are symmetric tropical Metzler matrices, we define the *tropical Metzler spectrahedron* associated with $Q^{(1)}, \dots, Q^{(n)}$ as the set of the points $x \in \mathbb{T}^n$ that fulfill the following two conditions:

- for all $i \in [m]$, $Q^+_{ii}(x) \geqslant Q^-_{ii}(x)$;
- for all $i, j \in [m]$, $i < j$, $Q^+_{ii}(x) \odot Q^+_{jj}(x) \geqslant (Q_{ij}(x))^{\odot 2}$.

Observe that the term $Q_{ij}(x)$ $(i \neq j)$ is well defined for any $x \in \mathbb{T}^n$ thanks to the Metzler property of the matrices $Q^{(k)}$.

In what follows, we denote by $\mathcal{S}$ the set described in Definition 3.3. With standard notation, the constraints defining this set respectively read: for all $i \in [m]$,

$$\max_{Q^{(k)}_{ii} \in \mathbb{T}_+} \left(Q^{(k)}_{ii} + x_k\right) \geqslant \max_{Q^{(l)}_{ii} \in \mathbb{T}_-} \left(|Q^{(l)}_{ii}| + x_l\right), \tag{5}$$

and for all $i, j \in [m]$ such that $i < j$,

$$\max_{Q^{(k)}_{ii} \in \mathbb{T}_+} \left(Q^{(k)}_{ii} + x_k\right) + \max_{Q^{(k')}_{jj} \in \mathbb{T}_+} \left(Q^{(k')}_{jj} + x_{k'}\right) \geqslant 2 \max_{l\in[n]} \left(|Q^{(l)}_{ij}| + x_l\right). \tag{6}$$

The next theorem ensures that $\mathcal{S}$ is indeed a tropical spectrahedron. To this purpose, we explicitly construct a spectrahedron $\boldsymbol{\mathcal{S}} \subset \mathbb{K}^n_{\geqslant 0}$ verifying $\mathsf{val}(\boldsymbol{\mathcal{S}}) = \mathcal{S}$.

Theorem 3.4. *The set $\mathcal{S}$ is a tropical spectrahedron.*

Proof. Let us define the matrices $\boldsymbol{Q}^{(1)}, \dots, \boldsymbol{Q}^{(n)} \in \mathbb{K}^{m\times m}$ as follows:

- if $Q^{(k)}_{ij} \in \mathbb{T}_-$, then we set $\boldsymbol{Q}^{(k)}_{ij} := -t^{|Q^{(k)}_{ij}|}$;
- if $Q^{(k)}_{ij} \in \mathbb{T}_+$ (which, under our assumptions, can happen only if $i = j$), then $\boldsymbol{Q}^{(k)}_{ij} := mnt^{Q^{(k)}_{ij}}$;
- if $Q^{(k)}_{ij} = -\infty$, then $\boldsymbol{Q}^{(k)}_{ij} := 0$.

Consider the spectrahedron $\boldsymbol{\mathcal{S}} := \{\boldsymbol{x} \in \mathbb{K}^n_{\geqslant 0} \colon \boldsymbol{Q}(\boldsymbol{x}) \succcurlyeq 0\}$. We claim that $\mathsf{val}(\boldsymbol{\mathcal{S}}) = \mathcal{S}$.

First, let us show that $\mathsf{val}(\boldsymbol{\mathcal{S}}) \subset \mathcal{S}$. Let $\boldsymbol{x} \in \boldsymbol{\mathcal{S}}$. The nonnegativity constraints on the minors of order 1 and 2 of $\boldsymbol{Q}(\boldsymbol{x})$ imply that $\boldsymbol{Q}^+_{ii}(\boldsymbol{x}) \geqslant \boldsymbol{Q}^-_{ii}(\boldsymbol{x})$ for all $i \in [m]$, and $\boldsymbol{Q}^+(\boldsymbol{x})_{ii}\boldsymbol{Q}^+(\boldsymbol{x})_{jj} \geqslant$

$(\boldsymbol{Q}(\boldsymbol{x})_{ij})^2$ for all $i \neq j$. By Lemma 3.1, for any point $\boldsymbol{x} \in \mathbb{K}^n_{\succcurlyeq 0}$ we have $\mathsf{val}(\boldsymbol{Q}^+_{ii}(\boldsymbol{x})) = Q^+_{ii}(x)$, where $x = \mathsf{val}(\boldsymbol{x})$. Similarly, $\mathsf{val}(\boldsymbol{Q}^-_{ii}(\boldsymbol{x})) = Q^-_{ii}(x)$. Finally, since the $Q^{(k)}$ are tropical Metzler matrices, we have $\mathsf{val}(\boldsymbol{Q}_{ij}(\boldsymbol{x})) = |Q_{ij}(x)|$ for $i \neq j$. Since the map val is order preserving over $\mathbb{K}_{\succcurlyeq 0}$, we deduce that $x \in \mathcal{S}$.

Now, let us prove the inclusion $\mathcal{S} \subset \mathsf{val}(\boldsymbol{\mathcal{S}})$. To this end, we introduce the following set

$$\boldsymbol{\mathcal{S}}' = \left\{ \boldsymbol{x} \in \mathbb{K}^n_{\succcurlyeq 0} \colon \forall i, \boldsymbol{Q}(\boldsymbol{x})_{ii} \succcurlyeq 0\,, \text{ and } \forall i \neq j, \boldsymbol{Q}(\boldsymbol{x})_{ii}\boldsymbol{Q}(\boldsymbol{x})_{jj} \succcurlyeq (m-1)^2(\boldsymbol{Q}(\boldsymbol{x})_{ij})^2 \right\}.$$

and we show that $\mathcal{S} \subset \mathsf{val}(\boldsymbol{\mathcal{S}}')$. By Lemma E.1, we have $\boldsymbol{\mathcal{S}}' \subset \boldsymbol{\mathcal{S}}$, and this will provide the expected result. Take any $x \in \mathcal{S}$ and its lift $\boldsymbol{x}_k = t^{x_k}$, with the convention that $t^{-\infty} = 0$. First, as noted in the previous paragraph, we have $\mathsf{val}(\boldsymbol{Q}^+_{ii}(\boldsymbol{x})) = Q^+_{ii}(x)$ and $\mathsf{val}(\boldsymbol{Q}^-_{ii}(\boldsymbol{x})) = Q^-_{ii}(x)$. We have chosen the matrices $\boldsymbol{Q}^{(k)}$ and the point $\boldsymbol{x}$ in such a way that

$$\boldsymbol{Q}^+_{ii}(\boldsymbol{x}) = \sum_{Q^{(k)}_{ii} \in \mathbb{T}^+} mnt^{Q^{(k)}_{ii} + x_k} \succcurlyeq mnt^{Q^+_{ii}(x)}\,. \tag{7}$$

Similarly, we have $\boldsymbol{Q}^-_{ii}(\boldsymbol{x}) = \sum_{Q^{(k)}_{ii} \in \mathbb{T}^-} t^{|Q^{(k)}_{ii}| + x_k} \preccurlyeq nt^{Q^-_{ii}(x)}$. Since $Q^+_{ii}(x) \succcurlyeq Q^-_{ii}(x)$, we deduce that $\boldsymbol{Q}^-_{ii}(\boldsymbol{x}) \preccurlyeq \frac{1}{m}\boldsymbol{Q}^+_{ii}(\boldsymbol{x})$, and so $\boldsymbol{Q}_{ii}(\boldsymbol{x}) \succcurlyeq (1 - \frac{1}{m})\boldsymbol{Q}^+_{ii}(\boldsymbol{x}) \succcurlyeq 0$. Second, for all $i \neq j$ we have $0 \succcurlyeq \boldsymbol{Q}_{ij}(\boldsymbol{x}) \succcurlyeq -nt^{|Q_{ij}(x)|}$. Using (7) and the fact that $Q^+_{ii}(x) \odot Q^+_{jj}(x) \succcurlyeq (Q_{ij}(x))^{\odot 2}$, we obtain $\boldsymbol{Q}_{ij}(\boldsymbol{x})^2 \preccurlyeq \frac{1}{m^2}\boldsymbol{Q}^+_{ii}(\boldsymbol{x})\boldsymbol{Q}^+_{jj}(\boldsymbol{x})$. Therefore, by the previous inequalities,

$$\boldsymbol{Q}_{ii}(\boldsymbol{x})\boldsymbol{Q}_{jj}(\boldsymbol{x}) - (m-1)^2\boldsymbol{Q}_{ij}(\boldsymbol{x})^2 \succcurlyeq (1 - \frac{1}{m})^2\boldsymbol{Q}^+_{ii}(\boldsymbol{x})\boldsymbol{Q}^+_{jj}(\boldsymbol{x}) - (m-1)^2\boldsymbol{Q}_{ij}(\boldsymbol{x})^2 \succcurlyeq 0\,.$$

Hence $\boldsymbol{x} \in \boldsymbol{\mathcal{S}}'$. □

3.2.2. *Effective tropicalization of nonarchimedean spectrahedra in the generic case.* Theorem 3.4 and its proof exhibits a construction of a lift $\boldsymbol{\mathcal{S}} = \{\boldsymbol{x} \in \mathbb{K}^n_{\succcurlyeq 0} \colon \boldsymbol{Q}(\boldsymbol{x}) \succcurlyeq 0\}$. This lift is obtained by taking matrices $\boldsymbol{Q}^{(k)}$ such that $\mathsf{sval}(\boldsymbol{Q}^{(k)}) = Q^{(k)}$. However, we have chosen the entries of these matrices in a very careful way to ensure that $\mathsf{val}(\boldsymbol{\mathcal{S}}) = \mathcal{S}$. It is natural to ask if the latter relation still holds for other choices of the coefficients $\boldsymbol{Q}^{(k)}_{ij}$. Said in a different way, we aim at identifying a larger class of spectrahedra $\boldsymbol{\mathcal{S}}$ whose tropicalization is described by the constraints (5) and (6) defining $\mathcal{S}$. In this section, we state that if the entries $Q^{(k)}_{ij}$ of the tropical matrices satisfy a certain genericity condition, then the tropicalization of *any* spectrahedron $\boldsymbol{\mathcal{S}} = \{\boldsymbol{x} \in \mathbb{K}^n_{\succcurlyeq 0} \colon \boldsymbol{Q}(\boldsymbol{x}) \succcurlyeq 0\}$ satisfying $\mathsf{sval}(\boldsymbol{Q}^{(k)}) = Q^{(k)}$ is given by the set $\mathcal{S}$. We start with a proposition which shows a weaker result on the tropicalization of the spectrahedron restricted to the open positive orthant $\mathbb{K}^n_{>0} := \{\boldsymbol{x} \in \mathbb{K}^n \colon \forall k, \boldsymbol{x}_k > 0\}$. Our genericity condition has two components: first, we assume that each matrix $Q^{(k)}$ has *nondegenerate tropical minors of order two*. In other words, we suppose that $Q^{(k)}_{ii} + Q^{(k)}_{jj} \neq 2|Q^{(k)}_{ij}|$ for every $i \neq j$ such that $Q^{(k)}_{ii}$ and $Q^{(k)}_{jj}$ belong to $\mathbb{T}_+$. This ensures that the closure of the interior of $\mathcal{S} \cap \mathbb{R}^n$ lies inside $\mathsf{val}(\boldsymbol{\mathcal{S}} \cap \mathbb{K}^n_{>0})$. Moreover, recall that a closed set $X \subset \mathbb{R}^n$ is called *regular* if $X = \mathrm{cl}(\mathrm{int}(X))$. Therefore, if we further assume that $\mathcal{S} \cap \mathbb{R}^n$ is regular, then we have the equality $\mathsf{val}(\boldsymbol{\mathcal{S}} \cap \mathbb{K}^n_{>0}) = \mathcal{S} \cap \mathbb{R}^n$.

Proposition 3.5. *Suppose that for every k the matrix $Q^{(k)}$ has nondegenerate tropical minors of order two and that the set $\mathcal{S} \cap \mathbb{R}^n$ is regular. Then $\mathsf{val}(\boldsymbol{\mathcal{S}} \cap \mathbb{K}^n_{>0}) = \mathcal{S} \cap \mathbb{R}^n$.*

We omit the proof of this statement, since it relies on techniques developed in a companion work on the tropicalization of semialgebraic sets.

Geometrically, the set $\mathcal{S} \cap \mathbb{R}^n$ is described by a polyhedral complex, i.e., a collection of convex polyhedra (called cells) such that the faces of every polyhedron belong to the collection, and the intersection of any two polyhedra is a face of both polyhedra. In our case, the maximal (inclusion-wise) cells are obtained by fixing, for each constraint of the form (5) and (6), some index k (and k') attaining the maximum. The tropical spectrahedron $\mathcal{S}$ is the union of these cells. The regularity condition amounts to the fact that all maximal cells are full-dimensional.

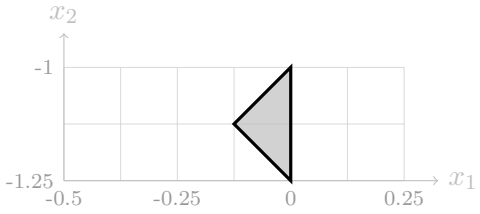

FIGURE 2. The tropicalization of the spectrahedron of Section 2.

The following lemma states that this condition is satisfied for any generic choice of the coefficients $Q_{ij}^{(k)}$:

Lemma 3.6. *Let $d = nm(m+1)/2$. There exists a finite set of hyperplanes in $\mathbb{R}^d$ such that, if the vector with entries $|Q_{ij}^{(k)}|$ (for $i \leqslant j$) is in $\mathbb{R}^d$ and belongs to none of these hyperplanes, then the assumptions of Proposition 3.5 are fulfilled.*

The *support* of a point $x \in \mathbb{T}^n$ is defined as the set of indices $k \in [n]$ such that $x_k \neq -\infty$. Given a nonempty subset $K \subset [n]$, and a set $X \subset \mathbb{T}^n$, we define the *stratum of X associated with K* as the subset of $\mathbb{R}^K$ formed by the projection $(x_k)_{k\in K}$ of the points $x \in X$ with support K. We deduce from Proposition 3.5 and Lemma 3.6:

Theorem 3.7. *If the coefficients $Q_{ij}^{(k)}$ are chosen generically, then all the strata of $\mathcal{S}$ are regular, and we have*

$$\mathsf{val}(\boldsymbol{\mathcal{S}}) = \mathcal{S}$$

for every spectrahedron $\boldsymbol{\mathcal{S}} = \{\boldsymbol{x} \in \mathbb{K}^n_{\succcurlyeq 0} \colon \boldsymbol{Q}(\boldsymbol{x}) \succcurlyeq 0\}$ such that $\mathsf{sval}(\boldsymbol{Q}^{(k)}) = Q^{(k)}$.

3.2.3. *Example.* Let us illustrate the results of this section on the spectrahedron defined in Section 2. The corresponding tropical matrices $Q^{(1)}, Q^{(2)}, Q^{(3)} \in \mathbb{T}^{3\times 3}$ are of Metzler type, and the associated tropical Metzler spectrahedron $\mathcal{S}$ is defined by the constraints:

$$\begin{aligned} \max\bigl(-1 + x_1, -5/4 + x_3\bigr) &\geqslant x_2 \\ \max\bigl(x_1 + x_3, -1/4 + 2x_3\bigr) &\geqslant 2x_1 \\ x_2 &\geqslant -7/4 + x_3 \\ \max\bigl(5/4 + x_1 + x_2, 1 + x_2 + x_3\bigr) &\geqslant 2x_3 \end{aligned}$$

The first inequality comes from (5) with $i = 2$, and the last three constraints from (6).[2] The intersection of $\mathcal{S}$ with the hyperplane $x_3 = 0$ is depicted on Figure 2. It can be checked that matrices $Q^{(k)}$ have nondegenerate tropical minors of order two and that all the strata of $\mathcal{S}$ are regular, so that the tropicalization of the spectrahedron of Section 2 is precisely described by the four inequalities given above.

4. Tropical spectrahedra and stochastic games

4.1. **Stochastic mean payoff games.** In this section, we present the class of games which is related to nonarchimedean semidefinite feasibility problems. For simplicity, we refer to them as stochastic mean payoff games, although as we shall see, this terminology usually corresponds to a larger class of games. This abuse of language is justified by the fact that the associated decision and computational problems are poly-time equivalent, as discussed below.

In our setting, a *stochastic mean payoff game* involves two players, Max and Min, which control disjoint sets of states. The states owned by Max and Min are respectively indexed by elements of $[m]$ and $[n]$. We will use the symbols i, j to refer to states of Player Max, and k, l to states of Player Min. Both players alternatively move a pawn over these states as follows.

[2]The constraints of the form (5) with $i = 1, 3$ are trivial.

When the pawn is on a state $k \in [n]$, Player Min chooses an action $a \in A^{(k)}$, which is defined as a subset of states of Max of cardinality 1 or 2: if $a = \{i\}$, then the pawn is moved to the state i, while if $a = \{i, j\}$ with $i \neq j$, it is moved to the state i (resp. j) with probability $1/2$. In both cases, Player Max receives from Player Min a reward denoted by r_k^a. Once the pawn is on a state $i \in [m]$, Player Max picks an action $b \in B^{(i)}$, where b is a subset of states of Min of cardinality 1. Then, Player Max moves the pawn to the state l such that $b = \{l\}$, and Player Min pays him a payment denoted by r_i^b.

We suppose that Player Min starts the game and that players can always make the next move, i.e., that $A^{(k)} \neq \emptyset$ and $B^{(i)} \neq \emptyset$ for all $i \in [m]$ and $k \in [n]$.

A *policy for Player Min* is a function mapping every state $k \in [n]$ to an action $\sigma(k)$ in $A^{(k)}$. Analogously, a *policy for Player Max* is a function τ such that $\tau(i) \in B^{(i)}$ for all $i \in [m]$. Suppose that the game starts from a state k^* of Player Min. When players play according to a couple (σ, τ) of policies, the movement of the pawn is described by a Markov chain. The average payoff of Player Max in the long-term is then defined as:

$$g_{k^*}(\sigma, \tau) = \lim_{N \to \infty} \mathbb{E}_{\sigma,\tau}\Big(\frac{1}{2N} \sum_{p=1}^{N} \big(r_{k_p}^{\sigma(k_p)} + r_{i_p}^{\tau(i_p)}\big)\Big) \ , \tag{8}$$

where the expectation $\mathbb{E}_{\sigma,\tau}$ is taken over all the trajectories $k_1, i_1, k_2, \ldots, i_p$ starting from $k_1 = k^*$ in the Markov chain.[3] The goal of Player Max is to find a policy which maximizes his average payoff, while Player Min aims at minimizing this quantity. The basic theorem of stochastic mean payoff games is the existence of "optimal" policies for both players. This was first proven by Liggett and Lippman [LL69].

Theorem 4.1. *There exists a pair of policies $(\overline{\sigma}, \overline{\tau})$ and a unique vector $\chi \in \mathbb{R}^n$ such that for all initial states $k \in [n]$, the following two conditions are satisfied:*

- *for each policy σ of Player Min, $\chi_k \leqslant g_k(\sigma, \overline{\tau})$;*
- *for each policy τ of Player Max, $\chi_k \geqslant g_k(\overline{\sigma}, \tau)$.*

The vector χ in Theorem 4.1 is referred to as the *value* of the game. Note that for every k, the quantity χ_k coincides with average payoff $g_k(\overline{\sigma}, \overline{\tau})$ associated with the couple of policies $(\overline{\sigma}, \overline{\tau})$. The first condition in Theorem 4.1 states that, by playing according to the policy $\overline{\tau}$, Player Max is certain to get an average payoff greater than or equal to the value χ_k associated with the initial state. Symmetrically, Player Min is ensured to limit her average loss to the quantity χ_k by following the policy $\overline{\sigma}$.

A state $k \in [n]$ is said to be *winning (for Player Max)* if the value χ_k of the game starting from the initial state k is nonnegative. We denote by Smpg the following decision problem: "given a stochastic mean payoff game, does there exist an initial state which is winning for Player Max?" It can be shown that if one of the players has no choice (i.e., has only one possible policy), then the problem of finding values and optimal policies can be solved in polynomial time by linear programming (see, e.g., [FV07, Section 2.9]). This readily implies that Smpg is in NP $\cap$ coNP (one guesses an optimal policy for one player, fixes it and then finds an optimal policy and the value of a 1-player game). However, as discussed in the introduction, the question whether there exists a poly-time algorithm to solve this problem is open.

Remark 4.2. In the literature, stochastic mean payoff games correspond to a larger class of problems [AM09]. These more general games admit a value and optimal policies as well. It turns out that the associated decision problem is poly-time equivalent to the simpler problem Smpg defined above. Moreover, these decision problems are poly-time equivalent to the problem of computing the value and a pair of optimal policies. We refer to Appendix C for a detailed proof.

[3] The limit in (8) can be shown to exist by the ergodic theorem, see Appendix A.1 for further details.

One of the possible approaches to analyze stochastic mean payoff games is to introduce the associated *Shapley operator*, which is a map F from $\mathbb{R}^n$ to itself defined as:

$$(F(x))_k = \min_{\substack{a\in A^{(k)}\\ a=\{i,j\}}} \Big(r_k^a + \frac{1}{2}\big(\max_{\substack{b\in B^{(i)}\\ b=\{l\}}} (r_i^b + x_l) + \max_{\substack{b\in B^{(j)}\\ b=\{l\}}} (r_j^b + x_l)\big)\Big)\,,$$

where we use the convention that $i = j$ when $a \in A^{(i)}$ is the singleton $\{i\}$. This map naturally extends to a self-map of $\mathbb{T}^n$, by setting $\frac{1}{2}(-\infty) = -\infty$.

The vectors x such that $x \leqslant F(x)$ may be thought of as nonlinear analogues of subharmonic functions. The existence of such a nontrivial vector x was shown in [AGG12] to be equivalent to the property that the mean payoff game has at least one winning state, in the case of deterministic games. The approach of [AGG12] is based in particular on a nonlinear fixed point theorem (Collatz-Wielandt theorem). The same arguments lead to the following result in the case of stochastic games.

Theorem 4.3. *We have the following properties:*

(i) Player Max has at least one winning initial state if and only if the set $\{x \in \mathbb{T}^n : x \leqslant F(x)\}$ *is nontrivial (i.e., contains a point different than* $(-\infty, \dots, -\infty)$*).*

(ii) Every initial state is winning for Max if and only if the set $\{x \in \mathbb{T}^n : x \leqslant F(x)\} \cap \mathbb{R}^n$ *is nonempty.*

For the convenience of the reader, the derivation of this result is detailed in Appendix B.

4.2. **Equivalence between stochastic games and tropical Metzler spectrahedra.** We now describe the connection between tropical spectrahedra and stochastic mean payoff games. Let us consider a tropical Metzler spectrahedron $\mathcal{S}$ associated with the matrices $Q^{(1)}, \dots, Q^{(n)} \in \mathbb{T}_{\pm}^{m\times m}$. We construct a stochastic mean payoff game Γ consisting of m states of Max and n states of Min. For each state $k \in [n]$ of Player Min, the set $A^{(k)}$ consists of:

- the actions $\{i\}$ with payment $-|Q_{ii}^{(k)}|$ for all $i \in [m]$ such that $Q_{ii}^{(k)} \in \mathbb{T}_-$;
- the actions $\{i, j\}$ with payment $-|Q_{ij}^{(k)}|$ for all $i < j$ such that $Q_{ij}^{(k)} \in \mathbb{T}_-$.

For every state $i \in [m]$ of Player Max, the set $B^{(i)}$ is formed by the actions $\{k\}$ with payment $Q_{ii}^{(k)}$ for all states $k \in [n]$ satisfying $Q_{ii}^{(k)} \in \mathbb{T}_+$.

Recall that in the games which we consider, every state has to be equipped with at least one action, i.e., the sets $A^{(k)}$ and $B^{(i)}$ must be nonempty. In consequence, our construction is valid provided that the following assumption on the matrices $Q^{(k)}$ is satisfied:

Assumption A. (a) For all $k \in [n]$, the matrix $Q^{(k)}$ has at least one coefficient in $\mathbb{T}_-$.
(b) For all $i \in [m]$, there exists $k \in [n]$ such that the diagonal coefficient $Q_{ii}^{(k)}$ belongs to $\mathbb{T}_+$.

Concerning the nontriviality of tropical spectrahedra, Assumption A can be made without loss of generality, up to extracting submatrices from $Q^{(k)}$ or eliminating some of them. We refer to Lemma E.2 for more details.

The relationship between the tropical spectrahedron $\mathcal{S}$ and the game Γ is given by the following lemma:

Lemma 4.4. *Let* $F : \mathbb{T}^n \to \mathbb{T}^n$ *be the Shapley operator associated with* Γ*. Then we have* $\mathcal{S} = \{x \in \mathbb{T}^n : x \leqslant F(x)\}$.

Proof. By construction of the game Γ, a vector $x \in \mathbb{T}^n$ satisfies $x \leqslant F(x)$ if and only if for all $k \in [n]$,

$$x_k \leqslant \min_{Q_{ij}^{(k)}\in\mathbb{T}_-} \Big(-|Q_{ij}^{(k)}| + \frac{1}{2}\big(\max_{Q_{ii}^{(l)}\in\mathbb{T}_+} (Q_{ii}^{(l)} + x_l) + \max_{Q_{jj}^{(l)}\in\mathbb{T}_+} (Q_{jj}^{(l)} + x_l)\big)\Big)\,,$$

or, equivalently, for all $i, j \in [m]$,

$$2\Big(\max_{Q_{ij}^{(k)}\in\mathbb{T}_-} \big(|Q_{ij}^{(k)}| + x_k\big)\Big) \leqslant \max_{Q_{ii}^{(l)}\in\mathbb{T}_+} \big(Q_{ii}^{(l)} + x_l\big) + \max_{Q_{jj}^{(l)}\in\mathbb{T}_+} \big(Q_{jj}^{(l)} + x_l\big)\,.$$

By distinguishing whether i and j are equal in these inequalities, and recalling that $Q_{ij}^{(k)}$ is in $\mathbb{T}_- \cup \{-\infty\}$ for all $k \in [n]$ when $i \neq j$, we recover the constraints given in (5) and (6). □

Then, from Theorem 4.3 and Lemma 4.4, we obtain:

Corollary 4.5. *The tropical spectrahedron $\mathcal{S}$ is nontrivial if and only if the stochastic game Γ has at least one winning initial state. Moreover, $\mathcal{S}$ contains a point $x \in \mathbb{R}^n$ if and only if all initial states of Γ are winning.*

Along the same lines, we can reciprocally associate a tropical spectrahedron with any stochastic game. In more details, starting from a stochastic game Γ specified as in Section 4.1, we construct symmetric Metzler matrices $Q^{(1)}, \dots, Q^{(n)} \in \mathbb{T}_{\pm}^{m \times m}$ from the actions of the players. We define $Q_{ij}^{(k)} = Q_{ji}^{(k)} := \ominus(-r_k^a)$ for all $k \in [n]$ and $i, j \in [m]$ such that $i \neq j$ and $a = \{i, j\}$ is an available action of Player Min at the state k. Diagonal coefficients of the matrices $Q^{(k)}$ are defined according to whether the actions $a := \{i\}$ and $b := \{k\}$ are available from the states k and i respectively:

- if $a \in A^{(k)}$ and $b \notin B^{(i)}$, we set $Q_{ii}^{(k)} := \ominus(-r_k^a)$, while if $b \in B^{(i)}$ and $a \notin A^{(k)}$, we define $Q_{ii}^{(k)} := r_i^b$;
- if both actions a and b occur simultaneously, we set $Q_{ii}^{(k)}$ to $\ominus(-r_k^a)$ if $(-r_k^a) > r_i^b$, and to r_i^b if $r_i^b \geqslant (-r_k^a)$.

Finally, all the other entries of the matrices $Q^{(k)}$ are set to $-\infty$. With this construction, it can be verified that Lemma 4.4 and, subsequently, Corollary 4.5 are still valid.

We denote by TMSDFP the tropical Metzler semidefinite feasibility problem: "given symmetric tropical Metzler matrices $Q^{(1)}, \dots, Q^{(n)} \in \mathbb{T}_{\pm}^{m \times m}$ satisfying Assumption A, is the associated tropical Metzler spectrahedron trivial?" As a consequence of Corollary 4.5, we obtain the equivalence between the two decision problems SMPG and TMSDFP. This equivalence can be refined thanks to the fact that SMPG restricted to games with payments equal to 0 or ± 1 is poly-time equivalent to SMPG for arbitrary payments, see Corollary C.5 in Appendix. This yields the following result:

Corollary 4.6. *The problems* SMPG *and* TMSDFP *are poly-time equivalent. Furthermore, if either of these problems can be solved in pseudopolynomial time, then both of them can be solved in polynomial time.*

4.3. **Algorithms for tropical Metzler semidefinite feasibility problem.** In this section, we discuss algorithms that can solve TMSDFP thanks to the equivalence with SMPG established in Corollary 4.6. Let L denote the maximal number of bits needed to encode an entry from matrices $Q^{(1)}, \dots, Q^{(n)}$. An exponential bound for TMSDFP is achieved by a naive algorithm that enumerates all policies of one player and uses linear programming to solve the remaining 1-player game.

Theorem 4.7. *There is an algorithm that solves* TMSDFP *in*

$$\min\{n^m, m^{2n}\}\mathrm{Poly}(m, n, L)$$

arithmetic operations.

We refer to Appendix E for the proof of the results in this section. The best currently known bound for TMSDFP can be derived from the randomized subexponential algorithm of Halman [Hal07]. The distinctive feature of this algorithm is that it works in *strongly* subexponential time, i.e., its arithmetic complexity does not depend on L.

Theorem 4.8. *There is a randomized algorithm that solves* TMSDFP *in expected*

$$\exp(O(\sqrt{(m+n)\log(m+n)}))$$

arithmetic operations.

```
1: procedure CheckFeasibility(Q^(1), ..., Q^(n), ε)
2:   ▷ Q^(1), ..., Q^(n) ∈ T_±^{m×m}, Metzler matrices satisfying Assumption A, ε > 0 a numerical precision.
3:   u := 0 ∈ R^n, v := 0 ∈ R^n
4:   while max_k u_k > −ε and min_k u_k < ε do
5:     v := max(v, u), u := F(u)
         ▷ F is the Shapley operator of the game Γ associated to Q^(1), ..., Q^(n). The operation max
       on vectors is understood entrywise.
6:   done
7:   if max_k u_k ⩽ −ε then
8:     The tropical spectrahedron is trivial.
9:   else
10:    The vector v belongs to the tropical spectrahedron.
11:  end
12: end
```

FIGURE 3. Checking the feasibility of a tropical semidefinite problem by value iteration.

We next illustrate our approach by benchmarks. We rely here on an algorithm with poorer theoretical bounds, but which is well adapted practically to some large scale instances: value iteration. This algorithm may be thought of as a nonlinear analogue of the power algorithm to compute the dominant eigenvalue of a matrix. Its advantage lies in scalability. This leads to a procedure called CheckFeasibility, provided in Figure 3 and which checks the triviality of a tropical spectrahedron. We next establish the validity of this algorithm in a special situation.

Proposition 4.9. *Assume that the matrices $Q^{(1)}, \dots, Q^{(n)}$ have finite entries. Then, the mean payoff of the associated game Γ is independent of the initial state. Suppose in addition that this mean payoff differs from 0. Then, for all choices of $\varepsilon > 0$, the procedure* CheckFeasibility *terminates and is correct.*

Note that the situation in which the mean payoff is zero is degenerate: then, an infinitesimal perturbation of the entries can make the spectrahedron trivial or nontrivial. Value iteration cannot naturally handle such degenerate situations, which can be solved by different methods like the policy iteration algorithm of [ACTDG13] (with an implementation in exact arithmetics). Every iteration (while loop) of the procedure CheckFeasibility takes a time $O(nm^2)$, which is *linear* in the size of the input. The number of iterations can only be bounded by an exponential in the size of the input. However, our benchmarks indicate that the algorithm is fast when the instance is "far from being degenerate".

Example 4.10. Procedure CheckFeasibility applied to the nonarchimedean spectrahedron of Section 2, with $\varepsilon = 10^{-8}$, terminates in 20 iterations. It returns a vector

$$v = (1.05612\dots, 0.0204082\dots, 1.12755\dots)$$

which lies in the interior of the tropical spectrahedron shown in Figure 2. The rounded vector $(1.06, 0.02, 1.13)$ is also in the interior, and $\boldsymbol{x} := (t^{1.06}, t^{0.02}, t^{1.13})$ fulfills $\boldsymbol{Q}(\boldsymbol{x}) \succcurlyeq 0$.

We report in Table 1 experimental results for different values of (n, m). We chose all the $|Q_{ij}^{(k)}|$, for $i \leqslant j$, to be independent random variables uniformly distributed on $[0, 1]$. Moreover, the diagonal coefficients $Q_{ii}^{(k)}$ were chosen to have a positive tropical sign (they belong to $\mathbb{T}_+$). We took $\varepsilon = 10^{-8}$ (the performance was similar for $\varepsilon = 10^{-6}$ or $\varepsilon = 10^{-10}$). Our experiments were obtained using a C program, distributed as an ancillary file attached to this arXiv preprint for reproducibility purposes.[4] This program was compiled under Linux with gcc -O3, and executed on a single core of an Intel(R) i7-4600U CPU at 2.10GHz with 16GB RAM. We report the average execution time over 10 samples for every value of (n, m). The number of iterations did not exceed 731 on this benchmark, and, for most (n, m), it was limited to a few units. Indeed,

[4]The ancillary file can be downloaded from `http://arxiv.org/src/1603.06916/anc`.

(n,m)	$(50,10)$	$(50,40)$	$(50,50)$	$(50,100)$	$(50,1000)$
time	0.000065	0.000049	0.000077	0.000279	0.026802
(n,m)	$(100,10)$	$(100,15)$	$(100,80)$	$(100,100)$	$(100,1000)$
time	0.000025	0.000270	0.000366	0.000656	0.053944
(n,m)	$(1000,10)$	$(1000,50)$	$(1000,100)$	$(1000,200)$	$(1000,500)$
time	0.000233	0.073544	0.015305	0.027762	0.148714
(n,m)	$(2000,10)$	$(2000,70)$	$(2000,100)$	$(10000,150)$	$(10000,400)$
time	0.000487	1.852221	0.087536	19.919844	2.309174

TABLE 1. Execution time (in sec.) of Procedure CHECKFEASIBILITY on random instances.

random instances exhibit experimentally a phase transition: for a given (n,m), the system is either feasible with overwhelming probability, or infeasible with overwhelming probability, unless (n,m) lies in a tiny region of the parameter space. Value iteration quickly decides feasibility, except in regions close to the phase transition. This explains why the execution time does not increase monotonically with (n,m) in our experiments (we included both easy and hard values of (n,m)).

Remark 4.11. By Theorem 3.7, the procedure CHECKFEASIBILITY allows us to verify the feasibility of a Metzler nonarchimedean spectrahedron if the *genericity* condition of this theorem is satisfied. However, even when this genericity condition is not satisfied, the same procedure can be shown to check correctly the feasibility of the nonarchimedean problem assuming that the value of the mean payoff game Γ *differs from zero*, see Proposition E.3.

5. Concluding remarks

In this paper, we have shown that under a genericity condition on the valuations, solving feasibility semidefinite problems over the field of Puiseux series, in which the input matrices have a Metzler sign pattern, reduces to a well studied class of zero-sum stochastic games. This leads both to complexity bounds and to algorithms capable experimentally to solve large scale nonarchimedean instances. The interest is also to relate two different problems which both have unsettled complexities. This is the first exposition of this approach. Let us now point some limitations of our current results, together with possible developments of this work.

First, the present approach can be extended to handle non-Metzler sign patterns. However, the description of tropical spectrahedra is more involved in this case.

Moreover, it would be interesting to provide a characterization of all (not necessarily generic) tropical spectrahedra. Such a characterization is known for tropical polyhedra, and is given by the tropical Minkowski-Weyl theorem (see [DY07, Prop. 2.1] and [GK11, Th. 2]). Unfortunately, the proof relies on the classical Minkowski-Weyl theorem, which has no known analogue for spectrahedra. Currently, we are unable to give a full characterization of tropical spectrahedra.

Finally, an interesting question is to use the present approach to deal with the real case, i.e., to see t as a deformation parameter and use the fact that the semidefinite feasibility problem is easily solvable for a large parameter t, to infer the solvability for smaller values of t. However, we expect exponential bounds on the bit size of t to guarantee that the deformed feasibility problem and the nonarchimedean feasibility problem have the same answer.

Acknowledgments

We thank the referees for their detailed comments.

References

[ABGJ15] X. Allamigeon, P. Benchimol, S. Gaubert, and M. Joswig. Tropicalizing the simplex algorithm. *SIAM Journal on Discrete Mathematics*, 29(2), 2015.

[ACS14] D. Auger, P. Coucheney, and Y. Strozecki. Finding optimal strategies of almost acyclic simple stochastic games. In *Proceedings of TAMC'14*. Springer, 2014.

[ACTDG13] M. Akian, J. Cochet-Terrasson, S. Detournay, and S. Gaubert. Solving multichain stochastic games with mean payoff by policy iteration. In *IEEE 52nd Annual Conference on Decision and Control (CDC)*, 2013.

[AGG12] M. Akian, S. Gaubert, and A. Guterman. Tropical polyhedra are equivalent to mean payoff games. *Int. J. Algebra Comput.*, 22(1), 2012.

[AGH15] M. Akian, S. Gaubert, and A. Hochart. Ergodicity conditions for zero-sum games. *Discrete and Continuous Dynamical Systems, series A*, 35(9), 2015.

[AGS16] X. Allamigeon, S. Gaubert, and M. Skomra. Solving generic nonarchimedean semidefinite programs using stochastic game algorithms. In *Proceedings of ISSAC'16*. ACM, 2016.

[AM09] D. Andersson and P.B. Miltersen. The complexity of solving stochastic games on graphs. In *Proceedings of ISAAC'09*, number 5878 in LNCS. Springer, 2009.

[Bih02] F. Bihan. Viro method for the construction of real complete intersections. *Adv. Math.*, 169(2), 2002.

[BPR13] G. Blekherman, P.A. Parrilo, and T.R. Rekha. *Semidefinite Optimization and Convex Algebraic Geometry*. MOS-SIAM Series on Optimization. SIAM, 2013.

[But10] P. Butkovič. *Max-linear systems: theory and algorithms*. Springer Monographs in Mathematics. Springer-Verlag London, Ltd., London, 2010.

[Chu67] K.L. Chung. *Markov Chains With Stationary Transition Probabilities*, volume 104 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1967.

[Con92] A. Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203–224, 1992.

[dKV15] E. de Klerk and F. Vallentin. On the Turing model complexity of interior point methods for semidefinite programming. arXiv:1507.03549, 2015.

[DY07] M. Develin and J. Yu. Tropical polytopes and cellular resolutions. *Experimental Mathematics*, 16(3), 2007.

[FV07] J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer, 2007.

[GK11] S. Gaubert and R.D. Katz. Minimal half-spaces and external representation of tropical polyhedra. *Journal of Algebraic Combinatorics*, 33(3):325–348, 2011.

[GLS93] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Algorithms and Combinatorics. Springer-Verlag, 1993.

[GM12] B. Gärtner and J. Matoušek. *Approximation algorithms and semidefinite programming*. Springer, 2012.

[Hal07] N. Halman. Simple stochastic games, parity games, mean payoff games and discounted payoff games are all LP-type problems. *Algorithmica*, 49(1):37–50, 2007.

[HNSED15] D. Henrion, S. Naldi, and M. Safey El Din. Exact algorithms for linear matrix inequalities. arXiv:1508.03715, 2015.

[KM03] S. Kwek and K. Mehlhorn. Optimal search for rationals. *Info. Proc. Letters*, 86(1):23–26, 2003.

[Koh80] E. Kohlberg. Invariant half-lines of nonexpansive piecewise-linear transformations. *Mathematics of Operations Research*, 5(3):366–372, 1980.

[LL69] T.M. Liggett and S.A. Lippman. Stochastic games with perfect information and time average payoff. *SIAM Review*, 11(4):604–607, 1969.

[Mar10] T. Markwig. A field of generalized Puiseux series for tropical geometry. *Rendiconti del Seminario Matematico, Università e Politecnico di Torino*, 68(1):79–92, 2010.

[MC11] D. Monniaux and P. Corbineau. On the generation of Positivstellensatz witnesses in degenerate cases. In *Proceedings of ITP'11*, pages 249–264, 2011.

[Meg89] N. Megiddo. On the complexity of linear programming. In *Advances in economic theory*, volume 12 of *Econom. Soc. Monogr.* CUP, 1989.

[MS15] D. Maclagan and B. Sturmfels. *Introduction to tropical geometry*, volume 161 of *Graduate Studies in Mathematics*. American Mathematical Society, 2015.

[NRS10] J. Nie, K. Ranestad, and B. Sturmfels. The algebraic degree of semidefinite programming. *Math. Program.*, 122(2, Ser. A):379–405, 2010.

[Put05] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics. Wiley-Interscience, 2005.

[Ram97] M.V. Ramana. An exact duality theory for semidefinite programming and its complexity implications. *Mathematical Programming*, 77(1):129–162, 1997.

[Sch13] C. Scheiderer. Sums of squares of polynomials with rational coefficients. *J. European Math. Soc.*, 2013.

[Ste10] S. A. Steinberg. *Lattice-ordered rings and modules*. Springer, New York, 2010.

[Stu94] Bernd Sturmfels. Viro's theorem for complete intersections. *Ann. Scuola Norm. Sup. Pisa Cl. Sci. (4)*, 21(3):377–386, 1994.

[Vir89] O. Ya. Viro. Real plane algebraic curves: constructions with controlled topology. *Algebra i Analiz*, 1(5), 1989.

[Yu15] J. Yu. Tropicalizing the positive semidefinite cone. *Proc. Amer. Math. Soc*, 143(5):1891–1895, 2015.

[ZP96] U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158(1–2):343–359, 1996.

Appendix A. Markov chains

A.1. **Characterization of payoff in stochastic mean payoff games.** In this section, we give the characterization of the payoff $g_k(\sigma,\tau)$ using the theory of Markov chains. Hereafter, we only consider Markov chains with finite state spaces.

Recall that a Markov chain can be described by a square *transition matrix* $P \in \mathbb{R}_{\geqslant 0}^{|E|\times|E|}$, where E is the set of states and, for all $(w,z) \in E$, P_{wz} denotes the probability that the Markov chain moves from state w to state z in one step. We say that $\pi \in \mathbb{R}_{\geqslant 0}^{|C|}$ is a *stationary distribution* on the set $C \subset E$ if $\pi_w = \sum_{z\in C} \pi_z P_{zw}$ for all $w \in C$ and $\sum_{w\in C} \pi_w = 1$. Furthermore, recall that a set $C \subset E$ is called a *recurrent class* if it has the following two properties: (a) if C is reached, then the chain will never leave it; (b) if C is reached, then every state in C will be visited infinitely many times (with probability one). Every Markov chain has at least one recurrent class, and every recurrent class has a unique stationary distribution. A state $w \in E$ is called *recurrent*, if it belongs to a recurrent class. Otherwise, it is called *transient*.

In order to give a characterization of payoff is stochastic games, fix a pair of policies (σ,τ). The 0-sum game obtained in this way can be regarded as a Markov chain on a space $E = [n] \uplus [m]$, where $\uplus$ denotes the disjoint union. Denote this chain by $(X_0, X_1, \dots)$ and let P be its transition matrix. Now, for any $w \in E$, let the random variable $T_w = \inf\{s \geqslant 1 \colon X_s = w\}$ denote the time of first return to w. By θ_w we denote the expected time of first return to w,

$$\theta_w = \mathbb{E}(T_w | X_0 = w)\,.$$

Furthermore, let u_w be defined as the payoff received by Player Max when the pawn leaves the state w. In other words, $u_w = r_w^{\sigma(w)}$ for $w \in [n]$ and $u_w = r_w^{\tau(w)}$ for $w \in [m]$. Let ξ_w be the expected payoff of Player Max before returning to w, i.e.,

$$\xi_w = \mathbb{E}\Big(\sum_{s=0}^{T_w - 1} u_{X_s} \Big| X_0 = w \Big)\,.$$

It can be proved (see Appendix A.3 below) that if $w \in E$ is a recurrent state, then θ_w and ξ_w are finite. The next theorem gives a characterization of the payoff $g_k(\sigma,\tau)$.

Theorem A.1. *If $k \in [n]$ is a fixed initial state, then the payoff $g_k(\sigma,\tau)$ is well defined and characterized as follows:*

(i) Suppose that k is recurrent in the 0-player game obtained by fixing (σ,τ). Let C denote the recurrent class to which k belongs. Then

$$g_k(\sigma,\tau) = \frac{\xi_k}{\theta_k} = \sum_{w\in C} u_w \pi_w\,,$$

where $(\pi_w)_{w\in C}$ is the stationary distribution on C. In particular, $g_k(\sigma,\tau)$ is constant for all initial states belonging to C. Furthermore, we have $\pi_w = \frac{1}{\theta_w}$ for all $w \in C$.

(ii) If k is transient and $C_1, \dots, C_p$ denote all the recurrent classes of 0-player game obtained by fixing (σ,τ), then $g_k(\sigma,\tau) = \sum_{s=1}^{p} g_{k_s}(\sigma,\tau)\psi_s$, where, for all s, ψ_s denotes the probability that the 0-player game starting from k reaches the recurrent class C_s, and $k_s \in C_s \cap [n]$ is an arbitrary state of C_s controlled by Player Min.

Theorem A.1 is well known, but, for the convenience of the reader, in Appendix A.3, we show how it can be derived from standard statements about Markov chains provided in the textbook of Chung [Chu67].

A.2. **Example from Section 2.** In this section, we revisit the example presented in Section 2, see Figure 1, and present the computation of payoff using Theorem A.1. As noted previously, both players in this example game have only two policies: at state $\boxed{2}$ Player Max can choose the action that goes to $\textcircled{1}$ or the action that goes to $\textcircled{3}$, whereas at state $\textcircled{3}$ Player Min can choose the action that goes to $\{\boxed{1},\boxed{3}\}$ or the action that goes to $\{\boxed{2},\boxed{3}\}$. Suppose that Player Max chooses the action that goes to $\textcircled{1}$ and that Player Min chooses the action that goes to $\{\boxed{1},\boxed{3}\}$. The Markov chain obtained in this way has the transition matrix of form

$$P = \begin{bmatrix} 0 & U \\ V & 0 \end{bmatrix},$$

where U describes the probabilities of transition from circle states to square states, and V describes the probabilities of transition from square states to circle states, i.e.,

$$U = \begin{array}{c} \\ \textcircled{1} \\ \textcircled{2} \\ \textcircled{3} \end{array} \begin{array}{c} \begin{array}{ccc} \boxed{1} & \boxed{2} & \boxed{3} \end{array} \\ \begin{bmatrix} 1/2 & 1/2 & 0 \\ 0 & 1 & 0 \\ 1/2 & 0 & 1/2 \end{bmatrix} \end{array}, \quad V = \begin{array}{c} \\ \boxed{1} \\ \boxed{2} \\ \boxed{3} \end{array} \begin{array}{c} \begin{array}{ccc} \textcircled{1} & \textcircled{2} & \textcircled{3} \end{array} \\ \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \end{array}.$$

This chain has only one recurrent class and all states belong to this class. Moreover, it is easy to verify that

$$\pi = \frac{1}{10}(2,1,2,2,2,1)$$

is the stationary distribution of this chain. Therefore, by Theorem A.1, the payoff of Player Max is equal to

$$g\Big(\{\boxed{1},\boxed{3}\},\{\textcircled{1}\}\Big) = \frac{1}{10}(-\frac{6}{4}+2-2+\frac{9}{4}) = \frac{3}{40}.$$

Similarly, if Player Max chooses the action that goes to $\textcircled{1}$ and Player Min chooses the action that goes to $\{\boxed{2},\boxed{3}\}$, then the transition matrix of the generated Markov chain has the form $P' = \left[\begin{smallmatrix} 0 & U' \\ V & 0 \end{smallmatrix}\right]$, where

$$U' = \begin{array}{c} \\ \textcircled{1} \\ \textcircled{2} \\ \textcircled{3} \end{array} \begin{array}{c} \begin{array}{ccc} \boxed{1} & \boxed{2} & \boxed{3} \end{array} \\ \begin{bmatrix} 1/2 & 1/2 & 0 \\ 0 & 1 & 0 \\ 0 & 1/2 & 1/2 \end{bmatrix} \end{array}$$

and V is the same as previously. In this case the chain also has only one recurrent class and every state belongs to this class. Furthermore,

$$\pi' = \frac{1}{14}(4,1,2,2,4,1)$$

is the stationary distribution of this chain. Hence the payoff of Player Max is equal to

$$g\Big(\{\boxed{2},\boxed{3}\},\{\textcircled{1}\}\Big) = \frac{1}{14}(2-4+\frac{9}{4}) = \frac{1}{56}.$$

In both cases, the payoff of Player Max is positive. Therefore, if Player Max chooses the action that goes to $\textcircled{1}$, then he is guaranteed to obtain a nonnegative payoff. One can check, by doing the same calculations for the remaining policies, that $\Big(\{\boxed{2},\boxed{3}\},\{\textcircled{1}\}\Big)$ is the unique couple of optimal policies in this game.

A.3. **A proof of Theorem A.1.** In this section, we present the derivation of Theorem A.1 based on the book of Chung [Chu67]. Let μ_{wz} we denote the probability that the Markov chain starting from $w \in E$ will reach $z \in E$ at least once, $\mu_{wz} = \mathbb{P}(\exists s \geqslant 1, X_s = z | X_0 = w)$. By definition, the state w is recurrent if $\mu_{ww} = 1$, and it is transient otherwise. By ζ_{wz} we denote the expected number of visits in z before returning to w, i.e.,

$$\zeta_{wz} = \mathbb{E}\Big(\sum_{s=0}^{T_w-1} \mathbf{1}_{\{X_s=z\}} \Big| X_0 = w\Big).$$

The following theorem describes the ergodic behavior of any finite (or countable) Markov chain.

Theorem A.2. *The Cesaro limit*

$$\lim_{N\to\infty} \frac{1}{N} \sum_{s=0}^{N} P^s$$

is well defined (we will denote it by M). Moreover, the entries of M are given as follows: if z is a transient state, then $M_{wz} = 0$ for all w. If z is recurrent, then $M_{wz} = \frac{\mu_{wz}}{\theta_z}$ for all w.

Proof. See [Chu67, Part I, §6, Theorem 4 and its Corollary]. □

Remark A.3. Note that the theorem above does not state that $\theta_z < \infty$ if the state z is recurrent. (We work under the convention that $\frac{a}{+\infty} = 0$ for all finite a.) Nevertheless, if the chain is finite, then we have $\theta_z < \infty$ for all recurrent states z, and this can be deduced as a corollary of the theorem above, as discussed below.

Observe that M is a stochastic matrix (as a limit of stochastic matrices). Moreover, we have $MP = PM = M$. This leads to the following corollary.

Corollary A.4. *If $C \subset E$ is a recurrent class, then $M_{wz} = \frac{1}{\theta_z}$ for all $w, z \in C$. Furthermore, $(\pi_w)_{w\in C}$ defined as $\pi_w = \frac{1}{\theta_w}$ is the unique stationary distribution on C. In particular, if $w \in E$ is a recurrent state, then $\theta_w < \infty$.*

Proof. Let first claim follows immediately from Theorem A.2. We will prove that $(\pi_w)_{w\in C}$ is a stationary distribution on C. Let P_C denote the square submatrix of P formed by the rows and columns of P with indices in C. We define M_C analogously. The first claim implies that M_C has identical rows. Since C is a recurrent class, we have $P_{wz'} = 0$ for all $w \in C, z' \notin C$. Hence, for all s we have $[P^s]_C = [P_C]^s$. Therefore $M_C = \lim_{N\to\infty} \frac{1}{N} \sum_{s=0}^{N} P_C^s$. Hence, M_C is stochastic. In other words, every row of M_C is a probability distribution on C and, since $M_C P_C = M_C$, this distribution is a stationary distribution on C. Since C is a recurrent class, stationary distribution on C has only strictly positive values. Hence we have $\theta_w < \infty$. The fact that the stationary distribution is unique follows from [Chu67, Part I, §7, Theorem 1]. □

The next theorem characterizes the relationship between entries of M and the values ζ_{wz}.

Theorem A.5. *If (w, z) belong to the same recurrent class, then $0 < \zeta_{wz} < \infty$. Moreover, if (f, g, h, w, z) are (not necessarily distinct) states belonging to the same recurrent class, then $M_{fg}/M_{hw} = \zeta_{zg}/\zeta_{zw}$.*

Proof. The fact that $0 < \zeta_{wz} < \infty$ follows from [Chu67, Part I, §9, Theorem 2 and 3]. Moreover, by Corollary A.4 we have $M_{wz} = \frac{1}{\theta_z} > 0$. Hence the claim follows from [Chu67, Part I, §9, Theorem 5 and remarks that precede it]. □

Corollary A.6. *If (w, z) belong to the same recurrent class, then*

$$M_{wz} = \frac{1}{\theta_z} = \frac{\zeta_{wz}}{\theta_w}\,.$$

Proof. By Theorem A.2 and Theorem A.5 we have $\theta_w/\theta_z = M_{zz}/M_{ww} = \zeta_{wz}/\zeta_{ww}$. By definition $\zeta_{ww} = 1$ and hence $\theta_w/\theta_z = \zeta_{wz}$. □

Proof of Theorem A.1. Let $k \in [n]$ be any initial state. Observe that for all $N \geqslant 1$ we have

$$\mathbb{E}\Big(\sum_{s=0}^{2N-1} u_{X_s} \Big| X_0 = k\Big) = [u + Pu + \cdots + P^{2N-1}u]_k\,.$$

Therefore $g_k(\sigma, \tau) = [Mu]_k$. In particular, $g_k(\sigma, \tau)$ is well defined.

Let us suppose that the initial state k is recurrent and denote its recurrent class by C. In this case, Corollary A.4 and Theorem A.2 imply that $g_k(\sigma,\tau) = \sum_{w\in C} u_w \pi_w$, where $(\pi_w)_{w\in C}$ is the stationary distribution on C. Moreover, Corollary A.6 gives the identity

$$\begin{aligned} g_k(\sigma,\tau) &= \frac{1}{\theta_k}\sum_{w\in C}\zeta_{kw}u_w \\ &= \frac{1}{\theta_k}\mathbb{E}\Big(\sum_{w\in C}\sum_{s=0}^{T_k-1} u_w \mathbf{1}_{\{X_s=w\}}\Big| X_0=k\Big) \\ &= \frac{1}{\theta_k}\mathbb{E}\Big(\sum_{s=0}^{T_k-1}\sum_{w\in E} u_w \mathbf{1}_{\{X_s=w\}}\Big| X_0=k\Big) \\ &= \frac{1}{\theta_k}\mathbb{E}\Big(\sum_{s=0}^{T_k-1} u_{X_s}\Big| X_0=k\Big) = \frac{\xi_k}{\theta_k}\,. \end{aligned}$$

Now, suppose that the initial state k is transient. Let $C_1,\dots,C_p$ denote the recurrent classes in our Markov chain. In this case, Theorem A.2 gives the identity

$$g_k(\sigma,\tau) = \sum_{s=1}^{p}\sum_{z\in C_s}\frac{\psi_s}{\theta_z}u_z = \sum_{s=1}^{p}\psi_s g_{k_s}(\sigma,\tau). \qquad \square$$

Remark A.7. We point out that if the strategies (σ,τ) are fixed, then the matrix M from Theorem A.2 can be computed using the algorithm presented in [Put05, Appendix A.3 and A.4]. This algorithm can be implemented to run in strongly polynomial complexity using the strongly polynomial version of Gaussian elimination (presented, for example, in [GLS93, Section 1.4]). Therefore, for fixed policies (σ,τ), the payoff $g(\sigma,\tau) = Mu$ can be computed in strongly polynomial complexity.

Appendix B. Stochastic games and Shapley operators

In this section, we study Shapley operators associated with stochastic mean payoff games and present the proof of Theorem 4.3. Before proceeding, let us introduce an additional definition. We say that a function $f\colon \mathbb{R}^n \to \mathbb{R}^n$ is *piecewise affine* if $\mathbb{R}^n$ can be partitioned into finitely many polyhedra such that f restricted to any of them is affine. We have the following easy lemmas:

Lemma B.1. *We have the following properties:*

(i) An affine combination of finitely many piecewise affine functions is piecewise affine.
(ii) The maximum (resp. minimum) of finitely many piecewise affine functions is piecewise affine.

Lemma B.2. *Suppose that functions $f^{(1)},\dots,f^{(p)}\colon \mathbb{R}^n \to \mathbb{R}^n$ are affine and define $f(x) = \max_s(f^{(s)}(x))$ (or $f(x) = \min_s(f^{(s)}(x))$). Let $x,y\in\mathbb{R}^n$ be any vectors and let $\gamma\in\mathbb{R}$. Then there exists $s^*\in[p]$ such that for all γ large enough we have $f(x+\gamma y) = f^{(s^*)}(x+\gamma y)$.*

In Section 4.1 we defined the Shapley operator of a stochastic mean payoff game as a function $F\colon \mathbb{T}^n \to \mathbb{T}^n$ of form

$$F(x)_k = \min_{\substack{a\in A^{(k)}\\ a=\{i,j\}}}\Big(r_k^a + \frac{1}{2}\big(\max_{\substack{b\in B^{(i)}\\ b=\{l\}}}(r_i^b + x_l) + \max_{\substack{b\in B^{(j)}\\ b=\{l\}}}(r_j^b + x_l)\big)\Big)\,,$$

where we use the convention that $(1/2)(-\infty) = -\infty$. Note that this operator can be written down in the following, more compact form:

$$F(x)_k = \min_{a\in A^{(k)}}\Big(r_k^a + \frac{1}{|a|}\sum_{j\in a}\max_{\substack{b\in B^{(j)}\\ b=\{l\}}}(r_j^b + x_l)\Big)\,.$$

Let us recall that $\mathbb{T}$ is endowed with the natural order $\leqslant$. This order generates an order topology on $\mathbb{T}$, i.e., the smallest topology that contains all intervals of form (a,b) and $[-\infty,b)$, where $a,b\in\mathbb{R}$. That also generates the product topology on $\mathbb{T}^n$. Is is easy to check that the functions

$$\begin{aligned}\mathbb{T}^2\ni(a,b)&\to a+b\in\mathbb{T},\\ \mathbb{T}^2\ni(a,b)&\to \max\{a,b\}\in\mathbb{T},\\ \mathbb{T}^2\ni(a,b)&\to \min\{a,b\}\in\mathbb{T},\\ \mathbb{T}\ni a&\to \frac{1}{2}a\in\mathbb{T}\end{aligned}$$

are continuous. If $x,y\in\mathbb{T}^n$, are vectors, then we denote $x\leqslant y$ if the inequality $x_k\leqslant y_k$ is fulfilled for every $k\in[n]$. The basic properties of Shapley operators are the following:

Lemma B.3. *The Shapley operator F has the following properties:*

(i) it is order preserving*, i.e., for all $x,y\in\mathbb{T}^n$ such that $x\leqslant y$ we have $F(x)\leqslant F(y)$;*
(ii) it is additively homogeneous*, i.e., if $e=(1,1,\dots,1)$ and $\lambda\in\mathbb{R}$, then $F(x+\lambda e)=F(x)+\lambda e$ for all $x\in\mathbb{T}^n$;*
(iii) F is continuous in the topology of $\mathbb{T}^n$;
(iv) $F_{|\mathbb{R}^n}$ is nonexpansive *in the supremum norm, i.e., for all $x,y\in\mathbb{R}^n$ we have $\|F(x)-F(y)\|\leqslant\|x-y\|$, where $\|x\|=\max_{k\in[n]}|x_k|$;*
(v) $F_{|\mathbb{R}^n}$ is piecewise affine.

Proof. The first two properties follow trivially from the definition of F. The third follows from the remarks about topology of $\mathbb{T}^n$ stated above. The fourth follows from the first two and the fact that F preserves $\mathbb{R}^n$. Indeed, for all $x,y\in\mathbb{R}^n$ we have $x\leqslant y+\|x-y\|e$. Therefore $F(x)\leqslant F(y+\|x-y\|e)=F(y)+\|x-y\|e$ and $F(x)-F(y)\leqslant\|x-y\|e$. Analogously, $F(y)-F(x)\leqslant\|x-y\|e$ and hence $\|F(x)-F(y)\|\leqslant\|x-y\|$. Finally, the last point follows from Lemma B.1. □

The next theorem was proven by Kohlberg.

Theorem B.4 ([Koh80])**.** *Suppose that function $f\colon\mathbb{R}^n\to\mathbb{R}^n$ is piecewise affine and nonexpansive in any norm. Then there exists a triple $(\eta,v,\gamma_0)\in\mathbb{R}^n\times\mathbb{R}^n\times\mathbb{R}$, $\gamma_0\geqslant 0$, such that $f(\eta+\gamma v)=\eta+(\gamma+1)v$ for all $\gamma\geqslant\gamma_0$. A pair (η,v) in called an* invariant half-line*. Furthermore, if (η_1,v_1) and (η_2,v_2) are invariant half-lines, then $v_1=v_2$.*

By Lemma B.3 and Theorem B.4, Shapley operator $F_{|\mathbb{R}^n}$ has an invariant half-line (η,v). This leads to the following corollary.

Corollary B.5. *For every $x\in\mathbb{R}^n$ we have*

$$\lim_{N\to\infty}\frac{1}{N}F^N(x)=v$$

(in particular, this limit exists).

Proof. Since (η,v) is an invariant half-line, we have

$$\lim_{N\to\infty}\frac{1}{N}F^N(\eta+\gamma v)=v$$

for every γ large enough. Now, take any $x\in\mathbb{R}^n$. Since F is nonexpansive in the supremum norm, we have $\|F^N(\eta+\gamma v)-F^N(x)\|\leqslant\|F^{N-1}(\eta+\gamma v)-F^{N-1}(x)\|\leqslant\ldots\leqslant\|\eta+\gamma v-x\|$. Therefore

$$\lim_{N\to\infty}\Big\|\frac{F^N(\eta+\gamma v)}{N}-\frac{F^N(x)}{N}\Big\|=0$$

and hence $\lim_{N\to\infty}\frac{1}{N}F^N(x)=v$. □

Theorem B.6 (Collatz-Wielandt property [AGG12, Lemma 2.8])**.** *Let $f\colon \mathbb{T}^n \to \mathbb{T}^n$ be an order-preserving, additively homogeneous, and continuous map. Then*

$$\lim_{N\to\infty} \max_{k\in[n]} \frac{f_k^N(x)}{N} = \max\{\lambda \in \mathbb{T}\colon \exists u \in \mathbb{T}^n, u \neq -\infty, f(u) \geqslant \lambda + u\}\,,$$

where we use the notation $\mathbb{T}^n \ni -\infty = (-\infty, -\infty, \dots, -\infty)$.

(Note that this Theorem is stated in [AGG12] with "sup" instead of "max" on the right hand side, but the proof presented in [AGG12] shows that this supremum is attained.) Again, by Lemma B.3, we may apply Theorem B.6 to Shapey operator F and its invariant half-line (η, v). Observe that, by Corollary B.5,

$$\lim_{N\to\infty} \max_{k\in[n]} \frac{F_k^N(x)}{N} = \max_{k\in[n]} v_k\,.$$

This leads to the following corollary.

Corollary B.7. *The set $\{x \in \mathbb{T}^n\colon x \leqslant F(x)\}$ contains a point different than $-\infty$ if and only if $\max_k v_k \geqslant 0$. It contains a real point if and only if $v_k \geqslant 0$ for all $k \in [n]$.*

Proof. By Corollary B.5 and Theorem B.6 we have

$$\max_k v_k = \max\{\lambda \in \mathbb{T}\colon \exists x \in \mathbb{T}^n, x \neq -\infty, F(x) \geqslant \lambda + x\}\,.$$

Hence $\max_k v_k \geqslant 0$ if and only if the set $\{x \in \mathbb{T}^n\colon x \leqslant F(x)\}$ contains a point different than $-\infty$. That proves the first statement.

To prove the second statement, suppose that $v_k \geqslant 0$ for all k. Then, for γ large enough, we have $F(\eta + \gamma v) = \eta + (\gamma + 1)v \geqslant \eta + \gamma v$, and $(\eta + \gamma v) \in \{x \in \mathbb{T}^n\colon x \leqslant F(x)\} \cap \mathbb{R}^n$. Conversely, take any $x \in \mathbb{R}^n$ such that $x \leqslant F(x)$. Since F is order preserving, we have $x \leqslant F^N(x)$ for all N. Hence $\frac{x}{N} \leqslant \frac{F^N(x)}{N}$ and, by Corollary B.5, $v_k \geqslant 0$ for all k. □

It remains to show that the vector v given by Theorem B.4 is equal to the doubled value of the game which generates the Shapley operator F. To do that, we need the following notation: for every fixed policy σ of Player Min, we denote by F^σ the Shapley operator of a 1-player game in which Player Min can only use σ. Analogously, for every fixed policy τ of Player Max, we denote by F^τ the Shapley operator of a 1-player game in which Player Min can only use τ. By $F^{\sigma,\tau}$ we denote the Shapley operator of a 0-player game in which Player Min can only use σ and Player Max can only use τ. In other words, these operators are defined as follows:

$$F^\sigma(x)_k = r_k^{\sigma(k)} + \frac{1}{|\sigma(k)|} \sum_{j\in\sigma(k)} \max_{\substack{b\in B^{(j)}\\ b=\{l\}}} (r_j^b + x_l)\,,$$

$$F^\tau(x)_k = \min_{a\in A^{(k)}} \Big(r_k^a + \frac{1}{|a|} \sum_{\substack{j\in a\\ \tau(j)=\{l\}}} (r_j^{\tau(j)} + x_l)\Big)\,,$$

$$F^{\sigma,\tau}(x)_k = r_k^{\sigma(k)} + \frac{1}{|\sigma(k)|} \sum_{\substack{j\in\sigma(k)\\ \tau(j)=\{l\}}} (r_j^{\tau(j)} + x_l)\,.$$

It is easy to verify that we have the following lemma:

Lemma B.8. *For every $x \in \mathbb{R}^n$ we have*

$$F(x) = \min_\sigma F^\sigma(x) = \max_\tau F^\tau(x)\,.$$

Furthermore, $F^\sigma(x) = \max_\tau F^{\sigma,\tau}(x)$, as well as $F^\tau(x) = \min_\sigma F^{\sigma,\tau}(x)$.

Moreover, an easy induction shows the following result:

Lemma B.9. *Fix an initial state $k = k_1 \in [n]$ and suppose that Player Min uses policy σ. Then the expected total payoff of Player Max after the N-th turn does not exceed $[F^\sigma]^N(0)_k$. In other words,*

$$\forall N\ \forall \tau,\ \mathbb{E}_{\sigma,\tau}\Big(\sum_{p=1}^{N}\big(r_{k_p}^{\sigma(k_p)} + r_{i_p}^{\tau(i_p)}\big)\Big) \leqslant [F^\sigma]^N(0)_k\,.$$

Analogously, if Player Max uses policy τ, then his expected total payoff after the N-th turn is not smaller than $[F^\tau]^N(0)_k$, i.e.,

$$\forall N\ \forall \sigma,\ \mathbb{E}_{\sigma,\tau}\Big(\sum_{p=1}^{N}\big(r_{k_p}^{\sigma(k_p)} + r_{i_p}^{\tau(i_p)}\big)\Big) \geqslant [F^\tau]^N(0)_k\,.$$

We are now ready to prove the existence of optimal policies in the stochastic mean payoff games. To do this, we need the following lemma.

Lemma B.10. *There exists a pair of policies $(\overline{\sigma},\overline{\tau})$ such that (η, v) is an invariant half-line of F, $F^{\overline{\sigma}}$, and $F^{\overline{\tau}}$.*

Proof. We start by showing the existence of $\overline{\sigma}$. Note that, by Lemma B.8, for any σ, $F^\sigma(\eta+\gamma v) = \max_\tau F^{\sigma,\tau}(\eta+\gamma v)$. Therefore, by Lemma B.2, there exists $\tau(\sigma)$ such that $F^\sigma(\eta+\gamma v) = F^{\sigma,\tau(\sigma)}(\eta+\gamma v)$ for all γ large enough. Take γ to be so large that this equality holds for all pairs $(\sigma,\tau(\sigma))$. By Lemma B.8, we have $F(\eta+\gamma v) = \max_\sigma F^\sigma(\eta+\gamma v) = \max_\sigma F^{\sigma,\tau(\sigma)}(\eta+\gamma v)$. We use Lemma B.2 again and obtain the existence of $\overline{\sigma}$ such that $F(\eta+\gamma v) = F^{\overline{\sigma},\tau(\overline{\sigma})}(\eta+\gamma v) = F^{\overline{\sigma}}(\eta+\gamma v)$ for all γ large enough. Therefore $\eta+(\gamma+1)v = F(\eta+\gamma v) = F^{\overline{\sigma}}(\eta+\gamma v)$ and (v,η) is an invariant half-line of $F^{\overline{\sigma}}$. The construction of $\overline{\tau}$ is analogous. □

Theorem B.11. *Policies $(\overline{\sigma},\overline{\tau})$ are optimal. Furthermore, the value of the game is equal to $v/2$.*

Proof. Suppose that Player Min uses $\overline{\sigma}$. By Lemma B.9, Lemma B.10, and Corollary B.5 we have

$$\forall \tau,\ g_k(\overline{\sigma},\tau) = \lim_{N\to\infty}\mathbb{E}_{\overline{\sigma},\tau}\Big(\frac{1}{2N}\sum_{p=1}^{N}\big(r_{k_p}^{\overline{\sigma}(k_p)} + r_{i_p}^{\tau(i_p)}\big)\Big) \leqslant \frac{1}{2}v_k.$$

Analogously, if Player Max uses $\overline{\tau}$, we have

$$\forall \sigma,\ g_k(\sigma,\overline{\tau}) = \lim_{N\to\infty}\mathbb{E}_{\sigma,\overline{\tau}}\Big(\frac{1}{2N}\sum_{p=1}^{N}\big(r_{k_p}^{\sigma(k_p)} + r_{i_p}^{\overline{\tau}(i_p)}\big)\Big) \geqslant \frac{1}{2}v_k\,.$$

Hence $(\overline{\sigma},\overline{\tau})$ are optimal and $g_k(\overline{\sigma},\overline{\tau}) = v_k/2$ for every $k \in [n]$. □

Finally, Theorem 4.3 follows directly from Theorem B.11 and Corollary B.7.

Appendix C. Equivalent forms of stochastic mean payoff games problem

In this section we present the (algorithmically) equivalent forms of stochastic mean payoff games, as mentioned in Remark 4.2. In order to do that, we need to introduce the notions of *simple stochastic games* and *stopping games*.

We start by defining the class of stopping games. We say that a pair of states $(i,k) \in [m]\times[n]$ is a *sink* if, when i or k is reached, the game loops forever between these two states. More formally, we have $A^{(k)} = \{\{i\}\}, B^{(i)} = \{\{k\}\}$. We say that the game is *stopping* if it has at least one sink and the probability that the game will reach a sink is equal one for every choice of policies (σ,τ) and every initial state. Note that if a game is stopping, then only the payoffs in sinks are important to determine its solution. Indeed, if we denote the sinks by $(i_1,k_1),\dots,(i_p,k_p)$ and the game is stopping, then the payoff of Player Max is given by

$$g_k(\sigma,\tau) = \sum_{s=1}^{p}(r_{k_s}^{\{i_s\}} + r_{i_i}^{\{k_s\}})\mu_s(k,\sigma,\tau),$$

where $\mu_s(k, \sigma, \tau)$ is the probability that the game starting from k reaches the sink (i_s, k_s) if the players use the policies (σ, τ). This expression depends only on payoffs in sinks.

Now, we introduce the class of simple stochastic games. We say that the stochastic game is *simple* if its set of states can be divided into three classes: states controlled by Player Min, Player Max, and Nature. Players Min and Max have only deterministic choices and Nature chooses the next state by tossing a coin. (To be coherent with the previous definition of stochastic game, we assume that Player Min controls the states of Nature — but she has no other choice than to toss a coin.) Formally, we suppose that for every $k \in [n]$ and every $a \in A^{(k)}$ we have the implication $|a| = 2 \implies |A^{(k)}| = 1$. Moreover, a simple game has two sinks - one with payoff 1 and the other with payoff 0. All other payoffs are equal to 0.

It may seem that solving simple games is indeed simpler that solving games in their full generality. Andersson and Miltersen [AM09] have shown that this is not the case. Let Smpg-comp denote the problem of finding the value and a pair of optimal policies in a stochastic mean payoff game.

Theorem C.1 ([AM09])**.** Smpg-comp *is poly-time equivalent to the problem of finding the values of stopping simple stochastic games.*

Note that this Theorem is stated in [AM09] without the word "stopping", but this is what the authors actually show. As already mentioned in Remark 4.2, this result is valid for a much wider class of games than those considered in this work. Furthermore, the authors of [AM09] define stopping simple stochastic games in a slightly more general way — in their version, the players can make multiple moves in a row. Nevertheless, observe that we can always add dummy states to a stopping simple stochastic game (i.e., states in which player has only one action) and obtain an equivalent game that belongs to the class considered here. This shows that, from the algorithmic point of view, these classes are equivalent.

Moreover, observe that if the game is both simple and stopping, then we can change its payoffs in sinks — instead of payoffs equal to 0 and 1, we can demand them to be equal to -1 and 1. This does not change the optimal policies of the game and acts as an affine transformation on the value vector. Hence, from now on we assume that payoffs in sinks of stopping simple stochastic games are equal to -1 and 1. We now show that the computational problem of finding values of simple games can be reduced to the decision problem. By $\textsc{Smpg}(k)$ we will denote the problem of deciding if a given state $k \in [n]$ is winning in the stochastic mean payoff game.

Lemma C.2. $\textsc{Smpg}(k)$ *restricted to stopping simple stochastic games is poly-time equivalent to* Smpg-comp*.*

Proof. It is obvious that $\textsc{Smpg}(k)$ can be reduced to Smpg-comp. We will show the opposite reduction. By Theorem C.1, Smpg-comp is poly-time reducible to the problem of finding values of stopping simple stochastic games. Fix such a game and let $\chi \in [-1, 1]^n$ denote its value.

First, we show an auxiliary reduction. Fix a rational number $\alpha \in [-1, 1] \cap \mathbb{Q}$ and an initial state $k \in [n]$. Suppose that we want to decide if $\chi_k \geqslant \alpha$. We will show that this is poly-time reducible (poly-time in the size of the game and the number of bits needed to encode α) to $\textsc{Smpg}(k)$. If $\alpha = -1$, then the answer is "yes". If $\alpha > -1$, then we may modify the game as follows: we suppose that when the sink with payoff 1 is reached, the game does not start to loop, but instead moves with probability $1 - \frac{1}{\alpha+1}$ to the sink with payoff -1 and with probability $\frac{1}{\alpha+1}$ to the (newly created) sink with payoff 1. Denote the value of the modified game by $\tilde{\chi}$. We have $\tilde{\chi}_k \geqslant 0 \iff \chi_k \geqslant \alpha$. The modified game is stopping but not simple. Nevertheless, we may apply the construct of Zwick and Paterson [ZP96, remarks preceding Theorem 6.1] and obtain (in polynomial time) a new game, which is stopping, simple, and has $\tilde{\chi}_k$ as the value at state k. This gives the auxiliary reduction.

Second, [ACS14, Lemma 1] (which improves [Con92, Lemma 2]) shows that χ_k is a rational number of form $a/b \in [-1, 1]$, where a, b are integers and $0 \leqslant b \leqslant 4^{n+1}$. Thus, χ_k can be found

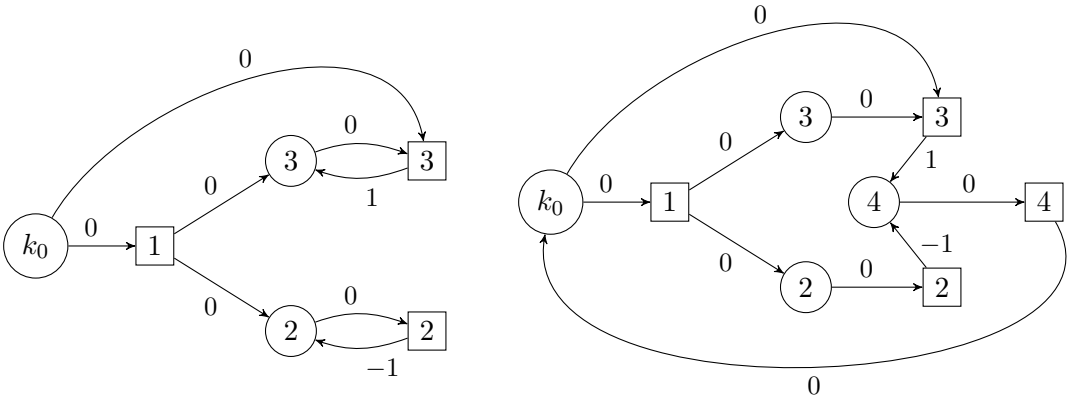

FIGURE 4. Transformation of stopping simple stochastic games. Circle states are controlled by Player Min. Numbers indicate payoffs received by Player Max after each move.

by the Kwek-Mehlhorn's algorithm (see [KM03]), using polynomially many queries to the oracle given by our auxiliary reduction. □

Finally, we want to show that SMPG(k) is poly-time reducible to SMPG. This requires an auxiliary construction which is presented on Figure 4. We take a stopping simple stochastic game Γ, fix an initial state $k_0 \in [n]$ and suppose that the sinks of Γ are indexed as $(n-1, m-1)$ (sink with payoff -1) and (n, m) (sink with payoff 1). Now, we modify the game as follows: we add two states, $n+1$ (controlled by Player Min) and $m+1$ (controlled by Player Max). At $n+1$, Player Min has only one possible action: to go to $m+1$; after this action Player Min pays 0 to Player Max. Moreover, at $m+1$ Player Max also has only one action: to go to k_0; after this action Player Max receives 0 from Player Min. Finally, we modify the sinks of Γ as follows: at $m-1$ (resp. m) Player Max has only one possible action: to go to $n+1$; after this action he receives -1 (resp. 1) from Player Min. Denote the modified game by $\overline{\Gamma}$. By construction, it is quite intuitive that the value of $\overline{\Gamma}$ does not depend on the initial state and that the state k_0 is winning in $\overline{\Gamma}$ if and only if it is is winning in Γ. To prove this formally, we use Theorem A.1. Henceforth, by $\Gamma^{\sigma,\tau}$ we denote the 0-player game obtained from Γ by fixing a pair of policies (σ, τ).

Corollary C.3. *The state k_0 is winning in Γ if and only if it is winning in $\overline{\Gamma}$. Moreover, the value of $\overline{\Gamma}$ is independent of the initial state.*

Proof. First, observe that there exists a natural bijection between policies of Γ and $\overline{\Gamma}$. Hence, we will use the same letters to denote policies in both games. Fix a pair of policies (σ, τ). Let $g(\sigma, \tau)$ (resp. $\bar{g}(\sigma, \tau)$) denote the payoff of Player Max in $\Gamma^{\sigma,\tau}$ (resp. in $\overline{\Gamma}^{\sigma,\tau}$). Since Γ was stopping, the 0-player game $\overline{\Gamma}^{\sigma,\tau}$ obtained from $\overline{\Gamma}$ by fixing the policies (σ, τ) has only one recurrent class and k_0 belongs to this class. By Theorem A.1 we see that $\bar{g}(\sigma, \tau)$ is constant for all initial states. Since (σ, τ) are arbitrary, this shows that the value of $\overline{\Gamma}$ does not depend on the choice of initial state. Furthermore, Theorem A.1 shows that $\bar{g}_{k_0}(\sigma, \tau) = g_{k_0}(\sigma, \tau)/\theta_{k_0}$, where $\theta_{k_0} = \theta_{k_0}(\sigma, \tau)$ is the expected time of first return to k_0.

Now, suppose that the value of Γ starting from k_0 is higher or equal than 0. Let τ^* denote the optimal policy of Player Max in Γ. For any policy σ of Player Min we have $g_{k_0}(\sigma, \tau^*) \geqslant 0$ and hence $\bar{g}_{k_0}(\sigma, \tau^*) \geqslant 0$. Therefore, τ^* is a winning (but not necessarily optimal) policy for Player Max in $\overline{\Gamma}$ starting from k_0. Thus, the value of $\overline{\Gamma}$ starting from k_0 is higher or equal than 0. The opposite implication is analogous. □

Corollary C.4. SMPG-COMP *is poly-time equivalent to* SMPG *restricted to games with payoffs in* $\{-1, 0, 1\}$.

Proof. By Lemma C.2, SMPG-COMP can be reduced to SMPG(k) restricted to stopping simple stochastic games. By the construction described above and Corollary C.3, SMPG(k) for stopping simple stochastic games can be reduced to SMPG restricted to games with payoffs in $\{-1, 0, 1\}$. The opposite reduction is trivial. □

Corollary C.5. SMPG *restricted to games with payoffs in* $\{-1, 0, 1\}$ *is poly-time equivalent to* SMPG *for general games.*

Proof. SMPG is trivially reducible to SMPG-COMP. Therefore, the claim follows from Corollary C.4. □

Appendix D. Analysis of the procedure CheckFeasibility

We now prove Proposition 4.9. To show that the mean payoff is independent of the initial state we rely on a result of [AGH15]. It suffices to check that the *recession function* of the Shapley operator, defined by

$$\hat{F}(x) := \lim_{\gamma\to\infty} \gamma^{-1} F(\gamma x)$$

has only trivial fixed points, i.e., that $\hat{F}(x) = x$ implies $x_1 = \cdots = x_n$. When all the entries $Q_{ij}^{(k)}$ are finite, one can verify that the k-th coordinate map of $\hat{F}$ is given by

$$\hat{F}_k(x) = \min_{i<j} \frac{1}{2}\big(\max_{Q_{ii}^{(l)}\in\mathbb{T}_+} x_l + \max_{Q_{jj}^{(l)}\in\mathbb{T}_+} x_l \big)\,,$$

and is therefore independent of the choice of $k \in [n]$. It follows that the only fixed points of $\hat{F}$ are constant vectors.

It will be convenient to use the following notation, for a vector $z \in \mathbb{R}^n$:

$$\mathbf{t}(z) := \max_i z_i, \qquad \mathbf{b}(z) := \min_i z_i,$$

so that the halting condition reads

$$\mathbf{t}(u) \leqslant -\varepsilon \text{ or } \mathbf{b}(u) \geqslant \varepsilon$$

The sequences $u^{(0)}, u^{(1)}, \dots$ and $v^{(0)}, v^{(1)}, \dots$ generated by the algorithm implemented in exact arithmetics satisfy, for $\ell \geqslant 1$,

$$u^{(\ell)} = F^\ell(0), \qquad v^{(\ell)} = \max(0, F(0), \dots, F^{\ell-1}(0)) \ .$$

We know from Theorem B.11 and Corollary B.5 that limit $\lim_\ell u^{(\ell)}/2\ell$ coincides with the mean payoff vector χ. The assumptions of the proposition imply that χ is a constant vector with nonzero entries, i.e., $\chi = (\lambda, \dots, \lambda)^\top$ where $\lambda \neq 0$. Therefore, $\lim_\ell \mathbf{b}(u^{(\ell)})/\ell = \lim_\ell \mathbf{t}(u^{(\ell)})/\ell = \lambda$. In particular, for ℓ sufficiently large, we have either $\mathbf{b}(u^{(\ell)}) > \varepsilon$, or $\mathbf{t}(u^{(\ell)}) < -\varepsilon$, depending on the sign of λ, and so the halting condition is ultimately satisfied. If $\mathbf{t}(u^{(\ell)}) < -\varepsilon$, then, $F^\ell(0) \leqslant -\varepsilon e$, where e is the unit vector. Since F is order preserving and commutes with the addition of a constant vector, we get, after an immediate induction, that $F^{\ell p}(0) \leqslant -p\varepsilon e$ holds for all p, and therefore, the mean payoff vector of the game, which coincides with $\lim_p F^{\ell p}/(2\ell p)$, has negative entries. By the Collatz-Wielandt property (Corollary B.7), this implies that the tropical spectrahedron $\{x : x \leqslant F(x)\}$ is reduced to the trivial vector with $-\infty$ entries.

If $\mathbf{t}(u^{(\ell)}) > \varepsilon$, then the dual argument shows that the mean payoff of the game $\lim_p F^{\ell p}/(2\ell p)$ has positive entries. Moreover, since F is order preserving, we have, thanks to the halting

condition $\mathbf{b}(F^\ell(0)) > \varepsilon$, and in particular, $F^\ell(0) \geqslant \varepsilon e$. Hence,

$$\begin{aligned} F(v) &= F(\max(0, F(0), \dots, F^{\ell-2}(0), F^{\ell-1}(0))) \\ &\geqslant \max(F(0), F^2(0), \dots, F^{\ell-1}(0), F^\ell(0)) \\ &\geqslant \max(F(0), F^2(0), \dots, F^{\ell-1}(0), \varepsilon e) \\ &\geqslant \max(F(0), F^2(0), \dots, F^{\ell-1}(0), 0) = v \end{aligned}$$

which implies that v belongs to the tropical spectrahedron. $\{x : x \leqslant F(x)\}$.

We note that the algorithm still terminates, and provides a correct yes/no answer, if the operator F is evaluated in fixed precision arithmetics, in such a way that u never differs from its true value by more than ε in the sup norm $\|\cdot\|$. Indeed, let $\tilde{u}^{(0)}, \tilde{u}^{(1)}, \dots$ denote the successive approximate values of the variables u which ae computed. Our assumption entails that $\|\tilde{u}^{(\ell)} - u^{(\ell)}\| \leqslant \varepsilon$ at every step ℓ of the algorithm. The termination proof relies on the fact that either $\mathbf{b}(u^{(\ell)})$ tends to infinity, or $\mathbf{t}(u^{(\ell)})$ tends to $-\infty$. Since $\|\tilde{u}^{(\ell)} - u^{(\ell)}\| \leqslant \varepsilon$, the analogous property is still satisfied by $\tilde{u}^{(\ell)}$, and so, the algorithm implemented with finite precision does terminate.

Now, if ℓ^* is the step at which the procedure terminates, we have either $\mathbf{b}(\tilde{u}^{(\ell^*)}) > \varepsilon$ or $\mathbf{t}(\tilde{u}^{(\ell^*)}) < -\varepsilon$, which entails that either $\mathbf{b}(u^{(\ell^*)}) > 0$ or $\mathbf{t}(u^{(\ell^*)}) < 0$. Reasoning as above, we deduce that the mean payoff of the game is positive in the first situation. and negative in the second one. Hence, the algorithm still decides correctly the triviality of the spectrahedron.

Appendix E. Other results

Lemma E.1. *Let $\boldsymbol{A} \in \mathbb{K}^{m\times m}$ be a symmetric matrix. Suppose that $\boldsymbol{A}$ has nonnegative entries on its diagonal and that the inequality $\boldsymbol{A}_{ii}\boldsymbol{A}_{jj} \geqslant (m-1)^2\boldsymbol{A}_{ij}^2$ holds for all pairs (i,j) such that $i \neq j$. Then $\boldsymbol{A}$ is positive semidefinite.*

Proof. If $\boldsymbol{A}$ is a zero matrix, then there is nothing to show. From now on we suppose that $\boldsymbol{A}$ has at least one nonzero entry. First, let us suppose that $\boldsymbol{A}$ has positive entries on its diagonal. In this case, let $\boldsymbol{B} \in \mathbb{K}^{m\times m}$, be the diagonal matrix which fulfills $\boldsymbol{B}_{ii} = \boldsymbol{A}_{ii}^{-1/2}$ for all i. Observe that $\boldsymbol{A}$ is positive semidefinite if and only if the matrix $\boldsymbol{D} = \boldsymbol{B}\boldsymbol{A}\boldsymbol{B}$ is positive semidefinite. Moreover, $\boldsymbol{D}$ has ones on its diagonal and $\boldsymbol{D}_{ij} = \boldsymbol{A}_{ii}^{-1/2}\boldsymbol{A}_{ij}\boldsymbol{A}_{jj}^{-1/2}$ for all $i \neq j$. Hence $|\boldsymbol{D}_{ij}| \leqslant 1/(m-1)$ for all $i \neq j$. Therefore $\boldsymbol{D}$ is diagonally dominant and hence positive semidefinite.

Now, if $\boldsymbol{A}$ has some zeros on its diagonal, let $I = \{i \in [m] \colon \boldsymbol{A}_{ii} \neq 0\}$. Since the inequality $\boldsymbol{A}_{ii}\boldsymbol{A}_{jj} \geqslant (m-1)^2\boldsymbol{A}_{ij}^2$ holds, we have $\boldsymbol{A}_{ij} = 0$ if either $i \notin I$ or $j \notin I$. Let $\boldsymbol{A}_I$ denote the submatrix formed by the rows and columns with indexes from I. Then $\boldsymbol{A}$ is positive semidefinite if and only if $\boldsymbol{A}_I$ is positive semidefinite. Finally, $\boldsymbol{A}_I$ is positive semidefinite by the considerations from the previous paragraph. □

Lemma E.2. *Let $Q^{(1)}, \dots, Q^{(n)} \in \mathbb{T}_\pm^{m\times m}$ be symmetric Metzler matrices, and $\mathcal{S}$ be the associated spectrahedron. We can build (in poly-time) symmetric Metzler matrices $R^{(1)}, \dots, R^{(q)} \in \mathbb{T}_\pm^{p\times p}$ ($p \leqslant m$, $q \leqslant n$) satisfying Assumption A such that the associated tropical spectrahedron is nontrivial if and only if $\mathcal{S}$ is nontrivial.*

Proof. We first examine Assumption A(a). Suppose that the matrix $Q^{(k)}$ has no coefficient in $\mathbb{T}_-$. In this case, the spectrahedron $\mathcal{S}$ is nontrivial, since it contains the vector $x \in \mathbb{T}^n$ such that $x_k = 0$ and $x_l = -\infty$ for all $l \neq k$.

Now, let us look at Assumption A(b). Consider $i \in [m]$ and suppose that no diagonal coefficient $Q_{ii}^{(k)}$ is in $\mathbb{T}_+$. We distinguish three cases:

- if the set K of indices $k \in [n]$ such that $Q_{ii}^{(k)} \in \mathbb{T}_-$ is nonempty, the relation $Q_{ii}^+(x) \geqslant Q_{ii}^-(x)$ enforces to have $x_k = -\infty$ for all $x \in \mathcal{S}$ and $k \in K$. This means that we can reduce the nontriviality of $\mathcal{S}$ to the nontriviality of the spectrahedron associated with the matrices $Q^{(k)}$ with $k \notin K$;

- if the aforementioned set K is empty, and the i-th row and column of the matrices $Q^{(k)}$ are all identically equal to $-\infty$, then we can remove all these rows and columns, and reduce to a problem with matrices of order $m-1$ over the variables $x_1, \dots, x_n$;
- if the set K is empty and some matrix $Q^{(k)}$ contains an entry different than $-\infty$ on its i-th row, namely $Q^{(k)}_{ij} \in \mathbb{T}_-$ with $i < j$, then the relation $Q^+_{ii}(x) \odot Q^+_{jj}(x) \geqslant (Q_{ij}(x))^{\odot 2}$ enforces $x_k = -\infty$ for any $x \in \mathcal{S}$. We consequently reduce the problem to the nontriviality of the spectrahedron associated with the matrices $Q^{(l)}$ with $l \neq k$. □

(Corollary 4.6). The fact that TMSDFP is poly-time equivalent to SMPG follows directly from Corollary 4.5. Moreover, the same statement shows that SMPG restricted to games with payoffs in $\{-1, 0, 1\}$ is poly-time equivalent to TMSDFP restricted to matrices with entries in

$$\{0, \ominus 0, \pm 1, \ominus(\pm 1), -\infty\}.$$

Therefore, by Corollary C.5, if either of these problems can be solved in pseudopolynomial time, then both are solvable in polynomial time. □

(Theorem 4.7). Take the matrices $Q^{(1)}, \dots, Q^{(n)}$ and consider the associated stochastic game Γ. If $m^{2n} \leqslant n^m$, then we do the following: for each policy σ of Player Min, we consider a 1-player game Γ^σ induced by fixing σ. The value of Γ^σ, denoted χ^σ, can then be found in polynomial time by linear programming (see [FV07, Section 2.9]). Moreover, for every fixed state $k \in [n]$, the value χ_k of Γ can be found by taking the minimum over σ of values χ^σ_k. Now, by Corollary 4.5, the tropical Metzler spectrahedron $\mathcal{S}$ associated with $Q^{(1)}, \dots, Q^{(n)}$ is nontrivial if and only if χ has at least one nonnegative entry. Since Player Min has at most m^{2n} policies in Γ, this procedure takes at most $m^{2n}\mathrm{Poly}(m, n, L)$ arithmetic operations. If $m^{2n} > n^m$, then we do the analogous operation for Player Max. □

(Theorem 4.8). As in the proof of Theorem 4.7, take the matrices $Q^{(1)}, \dots, Q^{(n)}$ and consider the associated game Γ. Now, we apply to Γ the first step of reduction presented in [AM09, Lemma 1 and Lemma 3] (i.e., the step presented on [AM09, Fig. 2]). This converts Γ in a strongly polynomial complexity to a game Γ_1, which has the same optimal policies as Γ and has the form considered by Halman [Hal07]. Moreover, Γ_1 has at most n min vertices and at most m max vertices (in Halman's terminology). Furthermore, Halman [Hal07, Theorem 4.1] gives an algorithm that finds optimal policies for games belonging to his class in the expected $\exp(O(\sqrt{p \log p}))$ arithmetic operations, where p is the total number of min and max vertices. Hence, a pair of optimal policies in Γ_1 (which is also optimal in Γ) can be found in the expected

$$\exp(O(\sqrt{(m+n)\log(m+n)}))$$

arithmetic operations. Once a pair of optimal policies is known, the value of Γ can be found in strongly polynomial complexity by Remark A.7. As in the proof of Theorem 4.7, we use the value of Γ to decide whether the tropical Metzler spectrahedron $\mathcal{S}$ is trivial. □

Proposition E.3. *Suppose that all the hypotheses of Proposition 4.9 are satisfied. Let $\mathcal{S}$ denote the tropical Metzler spectrahedron associated with $Q^{(1)}, \dots, Q^{(n)}$ and let $(\boldsymbol{Q}^{(k)})_k$ be any sequence of symmetric Metzler matrices such that for all $k \in [n]$ we have $\mathsf{sval}(\boldsymbol{Q}^{(k)}) = Q^{(k)}$. Denote $\boldsymbol{Q}(\boldsymbol{x}) = \boldsymbol{x}_1\boldsymbol{Q}_1 + \cdots + \boldsymbol{x}_n\boldsymbol{Q}_n$ and define $\boldsymbol{\mathcal{S}}$ as the nonarchimedean spectrahedron*

$$\boldsymbol{\mathcal{S}} = \{\boldsymbol{x} \in \mathbb{K}^n_{\geqslant 0} \colon \boldsymbol{Q}(\boldsymbol{x}) \succcurlyeq 0\}.$$

Then, $\boldsymbol{\mathcal{S}}$ is trivial if and only if the tropical spectrahedron $\mathcal{S}$ is trivial.

Proof. As in the proof of Theorem 3.4, we have the inclusion $\mathsf{val}(\boldsymbol{\mathcal{S}}) \subset \mathcal{S}$. This implies the "if" statement. To prove the opposite direction, assume that the tropical spectrahedron $\mathcal{S}$ is nontrivial. Then, the value of the game Γ must be nonnegative. Since we assumed in Proposition 4.9 that this value is not zero, it must be positive. Therefore, by the Collatz-Wielandt property (Theorem B.6), we can find a vector $x \in \mathbb{R}^n$ such that $(F(x))_k > x_k$ for all

$k \in [n]$ (here, F denotes the Shapley operator associated with Γ). Then, an inspection of the proof of Lemma 4.4 shows that for all $i, j \in [m]$ the strict inequality

$$2\Big(\max_{Q_{ij}^{(k)} \in \mathbb{T}_-} \big(|Q_{ij}^{(k)}| + x_k\big)\Big) < \max_{Q_{ii}^{(l)} \in \mathbb{T}_+} \big(Q_{ii}^{(l)} + x_l\big) + \max_{Q_{jj}^{(l)} \in \mathbb{T}_+} \big(Q_{jj}^{(l)} + x_l\big)$$

holds. Consider now the vector $\boldsymbol{x}^* := (t^{x_k})_{k\in[n]}$. The above strict inequalities imply that the matrix $\boldsymbol{Q}(\boldsymbol{x}^*)$ satisfies the conditions of Lemma E.1. Therefore, $\boldsymbol{x}^*$ belongs to $\boldsymbol{\mathcal{S}}$. □

INRIA and CMAP, École Polytechnique, CNRS, 91128 Palaiseau Cedex France
E-mail address: firstname.lastname@inria.fr